

# Are People With Cyber Security Training Worse at Checking Phishing Email Addresses? Testing The Automaticity Of Checking The Sender's Address.

Daniel Conway<sup>1</sup>[0009-0007-9963-9611], Marcus Butavicius<sup>2</sup> [0000-0003-0722-3912] , Kun Yu<sup>3</sup> [0000-0001-5138-6749] and Fang Chen<sup>4</sup> [0000-0003-4971-8729]

<sup>1,3,4</sup> University of Technology, Sydney, Australia  
(daniel.conway@student.uts.edu.au, kun.yu@uts.edu.au,  
fang.chen@uts.edu.au)

<sup>2</sup> Defence Science and Technology Group, Department of Defence, Australia (marcus.butavicius@defence.gov.au)

**Abstract.** Cyber security training emphasises checking the sender's email address to identify phishing emails. Dual process theories of cognition suggest that with practice such tactics can transition from effortful, analytic processes to involuntary, heuristics and become 'automatic'. We tested the automaticity of this email habit by developing a scale for cyber security experience and then deployed an interference task where participants (n=61) had to make a decision about text colour and ignore sender's addresses from either legitimate or phishing emails. A surprising result emerged: the more cyber security training participants had, the less interference they exhibited in the colour selection task and the more they were able to ignore the content of the sender's addresses. This suggests that evaluating sender's addresses does not fulfill the criterion for 'automatic' processes when practiced and that more experienced people seem to be able to ignore this important cue when extraneous task goals are present.

**Keywords:** Cyber Security, Human Factors, Decision Making, Dual Process Models, Automaticity, Phishing Attacks.

## 1 Introduction

Phishing emails remain the most common method of attack in a cyber-attacker's arsenal with globally an estimated 29 billion such emails being sent every day [19]. Successful attacks can have extremely deleterious effects on peoples' personal lives and is known to have cost industry in the U.S.A alone \$2.7 billion in 2018 [9].

While many Internet attacks can be thwarted via technical measures, as long as there is a human within a system who is capable of receiving emails - a single click on a malicious link can compromise that entire system [22]. This means that organisations attempting to reduce their attack surface must grapple with the challenges of 'human in the loop' systems in order to reduce their risk of breaches. This involves potential changes to the interface as well as looking at improvements in the training of employees

to increase their ability to detect phishing emails. A number of Human-Computer Interaction solutions have been fielded in response to this problem such as alerts that attempt to provide users information about possible attacks [27], however attackers have also used existing internet/software capabilities to increase the effectiveness of their attacks via approaches such as spoofing email addresses [11].

For this work we define phishing emails as generically constructed, mass-distributed emails with either an attachment containing malicious code, or a link to an external website designed to trick the reader into divulging information such as authorisation credentials [1]. These emails are constructed to appear like legitimate requests for information or a call to action that users must try to detect during their day-to-day use of their email systems [20].

Dual process models of cognition are based on a long-established body of evidence that there are two distinct systems involved in human decision-making [15]. System 1 thinking is characterised as fast, shallow, unconscious and effortless. It consists of the use of heuristics or rules-of-thumb such as the recognition heuristic, where we ascribe a higher value to something that we have encountered before and recognise – than something new [10]. The fact that this happens effectively instantaneously, with no effort, usually with little self-awareness and is extremely difficult to compensate for, exemplifies the automatic nature of this type of thought. The other process is characterised as slow, more effortful thought and has been dubbed System 2 thinking. This is suggested to be more exhaustive, rational and analytical but is less frequently engaged in since it is more cognitively demanding [2].

This human tendency to rely mostly on simple decision-making heuristics has important implications in the context of phishing emails. Dual process theories suggest that if users allow fast, automatic, and shallow processing to dominate their decision-making, they are more likely to fall victim to malicious emails [4,17,30]. In particular attackers now include cues that deliberately provoke a strong System 1 response in order to hinder deeper thought that may cast doubt on message source. However users engaging in System 2 processing to carefully evaluate the incoming message are more likely to recognise malicious intent and not respond to the attack [4,23,30].

For organisations with cyber security training regimes, a central focus is to help staff identify incoming phishing emails and one of the commonest diagnostic tools emphasised in these efforts is that of examining the sender's email address [3]. In this paper we examine this particular tactic of identifying suspicious emails and attempt to see if this method can be practiced to the point where it can be considered a System 1 process. Such research has implications not only for the theoretical application of dual process models in this area but also practical benefit in the design of cyber security awareness programs in that it would demonstrate that such practices can be not only transmitted but become 'second nature'. Furthermore understanding the automaticity of this highly diagnostic tactic has important ramifications for the interaction design of email systems.

## 2 Background

Within corporate cyber security training a central tactic that is taught to help users identify fraudulent senders is that of evaluating the sender's email address [3]. Attributes of sender's addresses that are taught as predicting malicious intent are: letter substitution (for example: noreply@mazon.com), unusual domain names (Eg: mailer.srvvscust-yaoelauzwb9446325@mntapjwaku.com), and domains that do not match the content of the email [3,14,21].

According to dual process theory, one of the primary characteristics of System 1 processes is that they are 'automatic' [29]. This implies that, once stimuli are perceived that match and trigger a System 1 process, they run to completion, impervious to efforts to stop them. Furthermore, dual process theories suggest that some processes that can transition from effortful, System 2 processes to automatic and effortless System 1 processes – through practice [25].

Based on the findings presented in Conway [6] we therefore hypothesised that repeated training may well result in users having practiced this particular process – of immediately looking at the sender's address to ascertain authenticity – to some form of automaticity – and potentially an effortless System 1 process. This current paper investigates whether a process such as looking for the sender's address can be practiced to the point of automaticity and whether such automaticity complies with the expectations that dual process theories attribute to a System 1 heuristic.

### 2.1 Assessing Experience

The experiment needed to measure individual differences in cyber security training and experience between participants. Accordingly, we developed a 'Cyber Security Training and Experience' measure using a small number of constructs that have been associated with cyber security experience in the literature.

Previous studies have shown e-mail knowledge and experience as protective against individual phishing victimization [8,14] while Halevi [12] found that users who were more aware of cyber risks will protect themselves more. As such we asked Q1: *"I have read about or heard about how to identify phishing emails."*

Training users to examine sender's email addresses and whether they match the content of the email is central to many existing corporate training platforms [3]. We expected this particular 'experience' to be one of the more powerful determinants of whether people had practiced the procedure of checking sender's addresses to automatic levels. Therefore we asked Q2: *"I have had one or more training sessions on how to identify phishing emails at work."*

However we were cognisant of the fact that people may have encountered formal training in situations other than at work (Jakobsson, 2007) and therefore wanted to capture this eventuality. Therefore we asked Q3: *"I have had one or more training sessions on how to identify phishing emails elsewhere (not at work)."*

Apart from awareness training in the forms of educational materials, many corporate entities now run regular phishing 'drills' with their staff where they (or a third party) send emails to staff (un-announced) emulating phishing attacks in order to assess

whether staff can identify them [7,16]. We therefore asked Q4: “*My company regularly sends out phishing drills / emails.*”

Self-efficacy has been shown to strongly relate to ability in terms of being cyber secure. In a Health Belief model based study, Ng (2009) found that self-efficacy (in this case of being able to adopt protective security measures) was the most important determinant of intention and behaviour of all the variables they examined. As such we asked Q5: “*I know how to identify phishing emails.*”

Previously being victimised has a fraught relationship with subsequent good cyber security practices. Downs [8] found that people who had previously been victimised did not consider the consequences of fraud differently to those who had not. Therefore Q6 asked: “*I have fallen victim to one or more phishing emails in the past.*”

And finally, there has been some preliminary evidence [5] that those who are confident and experienced with cyber security have a tendency to share their knowledge and teach others. We were therefore again interested if this particular self-efficacy based trait was associated with such a practice. Q7 asked: “*I teach others how to identify phishing emails.*”

## 2.2 Testing for automaticity

A varying number of attributes, ranging in number from 2 to 14 have been ascribed as to be necessary for automatic processing [26] but one of the most commonly referred to attributes is that of obligatoriness, or not being able to consciously control or stop the execution of such processes once triggered [26]. This attribute is therefore what we sought to test for in the experiment that follows.

We hypothesised that there would be an individual difference in people who have received extensive cyber security training, either formally or informally, in the way that they processed sender’s addresses, with those who have practiced the process of email identification to the point of automaticity being unable to disregard the content of a sender’s email address – even when instructed to do so. An experiment was devised, in many ways similar to the widely-used Stroop task [28], where the primary measure was the difference in reaction times between congruent and non-congruent trials that we will refer to as the ‘Phishing Email Sender’s Address Interference Task’.

The Stroop task typically presents two sources of information, one being the colour of the stimuli, and other being the text of stimuli. Participants are instructed to ignore the text presented and respond only to the colour. Trials where the colour of the text is the same as the text itself are considered congruent trials, and ones in which they are different are considered incongruent. The fundamental finding here is that regardless of instructions, people typically respond slower to incongruent cues. This suggests that regardless of attempts to attenuate attention to the information presented in the text, reading and then some level of subsequent conceptual priming is carried out automatically. The slower response for non-congruent cues is said to be the result of conflict between the two sources of information, where the precept of the colour extracted from the text causes interference with the response to the colour stimuli.

The experiment presented was similar in structure to the classic Stroop task in that participants were told that they were to make a decision based on the colour of two on-

screen elements and should respond with one button if these two cues were the same colour, and another button if they did not match. One of the two cues was a sender's email address, and the other a colour block. In order to strongly cue the precepts of phishing emails and encourage semantic processing of the text – we labeled the two response buttons 'Phishing email' and 'Legitimate email'. Participants were, however, told to ignore this labeling as well as the content of the sender's address and respond only in response to whether the colours of the two cues matched.

We therefore formulated trials as being either congruent or incongruent. Congruent trials were those where the label of the correct response button matched the status of sender's address regardless of stimuli colour, i.e., pressing the 'Phishing email' button when the sender's address was a phishing email or pressing the 'Legitimate email' button when the sender's address was a legitimate email. Incongruent responses were those where the correct button to be pressed (according to the colour matching task) did not match status of the sender's address, i.e.: pressing 'Phishing email' when the sender's address was legitimate or pressing 'Legitimate' email when the sender's address was a phishing email.

The experiment that follows therefore presented trials of the Stroop-like attention based task, as well as then assessing participants cyber security training level via the instrument presented above. In doing so we expected to see:

**Hypothesis 1:** Participants would exhibit slower reaction times for incongruent trials than for congruent trials.

**Hypothesis 2:** This effect would be greater in those participants with extensive cyber security training.

### 3 Method

A single task and a number of questions measuring individual differences were presented sequentially within the same session using an online experiment coded in PsychoPy [24].

#### 3.1 Participants

Participants were recruited via the web-based recruitment platform Prolific ([www.prolific.co](http://www.prolific.co)). 66 participants were selected with the following criterion: English speaking from USA, UK, Ireland, Australia, NZ or SA. 33 were fulltime employed in 'business management and administration' or 'finance' according to the user data provided by Prolific. Removing incomplete responses, 61 responses remained (W: 31, M: 30), age: ( $M = 31.14$ ,  $SD = 12.55$ ). Mean experiment duration was 9 minutes 36 seconds.

#### 3.2 Materials

The sender email addresses presented were taken from real emails, both phishing (e.g., [accounts@mazon.com](mailto:accounts@mazon.com)) and legitimate (e.g., [epetition@parliament.nsw.gov.au](mailto:epetition@parliament.nsw.gov.au)). A small collection was made of phishing emails received by the author and colleagues






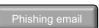


in the nine months before the experiment was began. These were then analysed as to characteristics noted by other authors [1,4,20] as indicative of phishing emails. We then selected emails that were most representative of these characteristics. As such we believe that the phishing email sender addresses presented were broadly representative of phishing emails in circulation at the time of writing. The legitimate email addresses presented were selected from legitimate emails similar in format to phishing emails in that they were not personal communications and contained a call to action.

### 3.3 Apparatus

The ‘Phishing email sender’s address congruency’ task consisted of trials where a sender’s email address and a solid block of colour were presented. Both address and colour block could be either brown or blue. Participants were instructed to press the left hand button when the colours did not match and the right hand button when they did. Of primary interest however was whether participants could ignore the semantic information contained in the email addresses. In order to provoke as much conflict as possible between the email content and the response, the left response button was labeled ‘Phishing email’ and the right ‘Legitimate email’. Participants were, however, instructed to ignore these labels and just respond to the colours of the two cues.

### 3.4 Dependent Measures

Reaction time was recorded for each trial along with the congruency of the colours presented and the response button pressed. Where the colours of the two cues matched and therefore the correct response was the ‘Legitimate email’ button, this was coded as ‘congruent’ when the content of the sender’s address was from a legitimate email and incongruent when from a phishing email. For trials where the colours were different and participants should have pressed the ‘Phishing email’ button, where the content of the sender’s address was a phishing email, trials were coded as congruent, and where the address was legitimate, coded as ‘incongruent’.

		Response coding used in analysis.	
		Congruent trials (Address content matches the button response)	Non-congruent trials (Address content does not match button response)
Primary task (colour congruency) not used in analysis	Colour cues congruent	Sender address: <a href="mailto:noreply@apple.com">noreply@apple.com</a> (legitimate) Colourblock:  Target button: 	Sender address: <a href="mailto:noreply@apple.com">noreply@apple.com</a> (legitimate) Colourblock:  Target button: 
	Colour cues non- congruent	Sender address: <a href="mailto:accounts@mazon.com">accounts@mazon.com</a> (phishing) Colourblock:  Target button: 	Sender address: <a href="mailto:accounts@mazon.com">accounts@mazon.com</a> (phishing) Colourblock:  Target button: 

**Fig. 1.** Stimuli and target response combination examples (not all combinations shown).

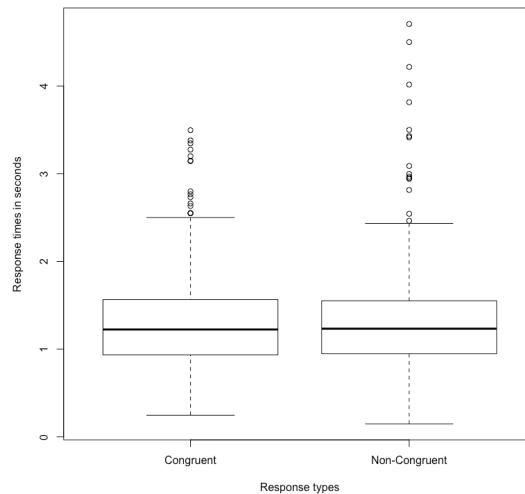
At the conclusion of the email/colour trials, participants were asked a series of questions to quantify their experience with or training in cyber security via the seven questions of the Cyber Security Training and Experience measure. All questions were provided as how much they agreed with each statement, with responses gathered via a five point Likert scale with all anchors labeled ('Strongly disagree', 'Disagree', 'Neutral', 'Agree', 'Strongly agree'). See table 1 for the text of all the questions. The resulting one to five scores for 7 questions were summed resulting in a total experience score with a range of 7 to 35.

### 3.5 Procedure

After gathering basic demographic information (age, gender and English speaking ability) instructions were provided as to the nature of the main task as well as two practice trials with explicit feedback. This was followed by all 16 trials of the Phishing Email Sender's Address Interference Task which were presented in a different random order for each participant, with eight consisting of a sender's address from a phishing email, and eight from legitimate emails. The colour congruency of the two cues were also distributed randomly throughout the trials. Then finally, the seven Cyber Security Training and Experience survey items were presented.

## 4 Results

Trials where response time was longer than five seconds (five trials) were removed as outliers (double the traditional  $1.5 \times \text{IQR}$  in order to remove less trials) leaving in 487 congruent trials and 484 incongruent trials.



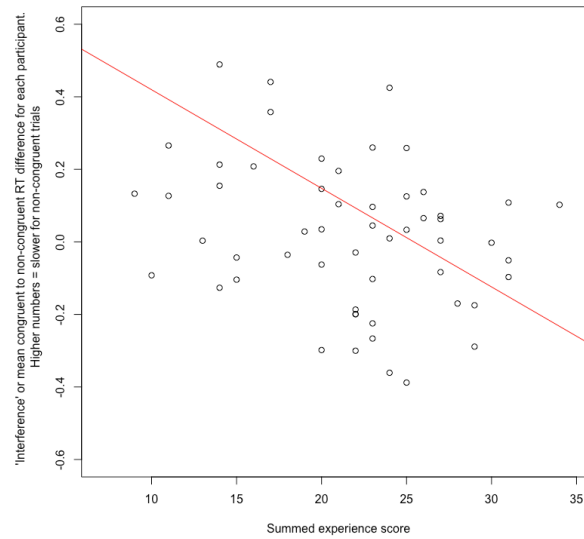
**Fig. 2.** Response times by whether the button pressed label was congruent to the status of the sender's address.

A one-sided Welch's t-test of reaction times for all congruent and non-congruent trials revealed that the mean reaction times of non-congruent trials,  $M = 1.327$  seconds, was not significantly higher than the mean reaction times of congruent trials,  $M = 1.305$  seconds,  $t = -0.635$ ,  $p = .262$ , (see figure 2).

#### 4.1 Interference Score

A total interference score was derived for each participant. This score was calculated by subtracting each participants' mean reaction time for non-congruent trials from their mean reaction time for congruent trials ( $M = 0.105$ ,  $SD = 0.47$ ). Positive values on this score indicated that the participant was slower for non-congruent trials and the magnitude of the score indicated the size of this effect. This interference score is used for much of the analysis that follows.

To examine the effect of experience score on interference a non-parametric, Rank Based regression analysis was carried out between participants' interference score and total experience scores. Simple rank-based linear regression indicated that for every increase of 1 in experience scores participants' interference score was 0.012 lower,  $\beta = -0.012$ ,  $p = .043$ . The experience score explained a statistically significant proportion of the interference score,  $R^2_{(\text{multiple robust})} = 0.08$ , although only accounted for 8% of the variation. See figure 3.



**Fig. 3.** Summed, total experience score by interference, for each participant

Since the individual difference questions included a neutral response and we were interested in the difference between those responded 'yes' and 'no' to these questions for subsequent analysis, we discarded neutral responses and divided the responding answers into two bins where answers of 1 ('strongly disagree') and 2 ('disagree') were

classified as ‘no’ answers and responses of 4 (‘agree’) and 5 (‘strongly agree’) were classified as ‘yes’ answers.

**Table 1:** Results of one-sided Welch’s t-tests and effect sizes for all seven experience questions.

Ques- tion Num- ber	Item	Mean Interfer- ence ‘No’ re- sponses	<i>n</i> of ‘No’ re- sponses	Mean Interfer- ence ‘Yes’ re- sponses	<i>n</i> of ‘Yes’ re- sponses	T-score	P value <sup>a</sup>	Cohen’s d	Effect size
1	<i>I have read about or heard about how to identify phishing emails.</i>	0.433	9	0.007	46	1.474	.177	.678	Large
2	<i>I have had one or more training sessions on how to identify phishing emails at work.</i>	0.35	22	-0.041	32	2.618	.015	.781	Large
3	<i>I have had one or more training sessions on how to identify phishing emails elsewhere (not at work).</i>	0.162	33	-0.073	19	2.484	.017	.631	Large
4	<i>My company regularly sends out phishing drills / emails.</i>	0.174	34	0.014	22	1.424	.162	.356	Medium
5	<i>I know how to identify phishing emails.</i>	0.684	8	0.0169	48	1.797	.114	.883	Large
6	<i>I have fallen victim to one or more phishing emails in the past.</i>	0.32	46	-0.064	12	1.337	.187	.321	Medium
7	<i>I teach others how to identify phishing emails.</i>	0.197	16	-0.062	35	2.404	.02	.603	Large

<sup>a</sup> In order to control the family-wise error rate when conducting multiple statistical tests, we used the Bonferroni adjustment for *p* values, resulting in a significance threshold of  $\alpha > .007$  – which none of tests achieved.

## 5 Discussion

We hypothesised that examining sender's address to identify malicious emails has been practiced to the point of automaticity in people with extensive cyber security training. We tested this automaticity via one of its the most commonly discussed attributes, namely the attribute of 'obligatoriness' or the inability to ignore the content of sender's address when presented with a task not requiring the processing of this information. However, our results did not support our first hypothesis, i.e., we found no evidence that responding to email addresses was automatic in our Phishing Email Sender's Address Interference Task. Specifically, there was no difference in response times between trials where the sender's address status matched the required response button (phishing / legitimate email) and when it did not match the response button. This implies that there was, overall, no cognitive interference between the content of the sender's email address and the required response suggesting that participants were able to ignore this information.

However, further analysis revealed a surprising phenomenon. Total Cyber Security Training and Experience scores, exhibited a weak, but significant correlation with the amount of interference that non-congruent trials presented. However, this correlation was in the opposite direction of that expected under our second hypothesis and instead participants with higher self reported experience were in fact better able to suppress the sender's address information when responding to the primary task.

A number of observations can be made to support this conclusion. Firstly, regardless of statistical significance, all seven questions trended in the same direction, suggesting that the main correlation presented is not the result of chance, but some kind of phenomena related to cyber experience. Secondly, the relative magnitudes of the effect sizes for each question correlate well with previous research. Here, questions with less stringent requirements such as Q1 (*'I have read about...'*) show a smaller effect size than does a question more likely to be more predictive of cyber training (Q2, *'I have received training...'*). And thirdly, questions involving constructs where previous research has shown a strong relationship with cyber security competence, such as self efficacy [13,18] demonstrated a larger effect size than those where the variable has been shown to have less effect, such as previous victimisation [8]. All of this taken together provides evidence that the phenomena described is closely related to cyber experience.

Our finding that those who are more experienced with cyber security are more adept at controlling their attention and ignoring task-irrelevant stimuli – even when this stimuli is strongly cued to interfere with the primary task – was counter to our predictions based on dual process theories of cognition. One possible explanation of this is that perhaps those with good attentional control are more likely to succeed in corporate environments and therefore receive more frequent and better quality cyber security training. However, this would mean that corporate training needs to take this tendency into account. Alternatively, those who have had more cyber security training and experience are also likely to have also had more experience in processing emails in a work environment. Such experience might have led to an increased ability to suppress detailed processing of email addresses when such processing is not relevant to the task

requirements such as quickly understanding and responding to emails within the context of ongoing email conversations. Again, this has implications for current training regimes. It is clear that more research is required to examine exactly how sender email addresses are processed in email tasks using a variety of experimental paradigms including, but not limited to, dual-task and eye-tracking experiments.

The implications of this finding are manifold. If some aspect of cyber security training results in people who are more able to ignore what they consider task-irrelevant stimuli, it suggests that this tendency would lead to increased victimisation. Within the HCI context, this suggests that new methods of providing visual emphasis to the Sender's address field such as highlighting may be required to influence users to better identify phishing emails targeting organisations.

## References

1. Nurul Akbar. 2014. Analysing Persuasion Principles in Phishing Emails. Retrieved May 29, 2021 from <http://essay.utwente.nl/66177/>
2. Lisa Feldman Barrett, Michele M. Tugade, and Randall W. Engle. 2004. Individual Differences in Working Memory Capacity and Dual-Process Theories of the Mind. *Psychological Bulletin* 130, 4: 553–573. <http://dx.doi.org/wwwproxy1.library.unsw.edu.au/10.1037/0033-2909.130.4.553>
3. Ladislav Burita, Ivo Klaban, and Tomas Racil. 2022. Education and Training Against Threat of Phishing Emails. *International Conference on Cyber Warfare and Security* 17, 1: 7–18. <https://doi.org/10.34190/iccws.17.1.28>
4. Marcus Butavicius, Ronnie Taib, and Simon J. Han. 2022. Why people keep falling for phishing scams: The effects of time pressure and deception cues on the detection of phishing emails. *Computers & Security*: 102937. <https://doi.org/10.1016/j.cose.2022.102937>
5. Dan Conway, Ronnie Taib, Mitch Harris, Kun Yu, Shlomo Berkovsky, and Fang Chen. 2017. A Qualitative Investigation of Bank Employee Experiences of Information Security and Phishing. 115–129. Retrieved March 8, 2021 from <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/conway>
6. Dan Conway, Kun Yu, Marcus Butavicius, and Fang Chen. 2022. Are phishing emails conflict problems? Dual process theory applied to an email identification task. *Manuscript in preparation*.
7. Ronald Dodge, Kathryn Coronges, and Ericka Rovira. 2012. Empirical Benefits of Training to Phishing Susceptibility. In *Information Security and Privacy Research*, 457–464. [https://doi.org/10.1007/978-3-642-30436-1\\_37](https://doi.org/10.1007/978-3-642-30436-1_37)
8. Julie S. Downs, Mandy B. Holbrook, and Lorrie Faith Cranor. 2006. Decision strategies and susceptibility to phishing. In *Proceedings of the second symposium on Usable privacy and security* (SOUPS '06), 79–90. <https://doi.org/10.1145/1143120.1143131>

9. FBI. *Internet Crime Complaint Center(IC3) | Annual Report 2018*. Federal Bureau of Investigations. Retrieved January 14, 2023 from <https://www.ic3.gov/Home/AnnualReports>
10. Gerd Gigerenzer and Daniel G. Goldstein. 2011. The recognition heuristic: A decade of research. *Judgment and Decision Making* 6, 1: 100–121.
11. Ashish Gupta, Ramesh Sharda, and Robert A. Greve. 2011. You’ve got email! Does it really matter to process emails now or later? *Information Systems Frontiers* 13, 5: 637–653. <https://doi.org/10.1007/s10796-010-9242-4>
12. Tzipora Halevi, Nasir Memon, and Oded Nov. 2015. *Spear-Phishing in the Wild: A Real-World Study of Personality, Phishing Self-Efficacy and Vulnerability to Spear-Phishing Attacks*. Social Science Research Network, Rochester, NY. <https://doi.org/10.2139/ssrn.2544742>
13. Princely Ifinedo. 2012. Understanding information systems security policy compliance: An integration of the theory of planned behavior and the protection motivation theory. *Computers & Security* 31, 1: 83–95. <https://doi.org/10.1016/j.cose.2011.10.007>
14. Markus Jakobsson. 2007. The human factor in phishing. *Privacy & Security of Consumer Information*.
15. Daniel Kahneman. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York, NY, US.
16. Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Protecting people from phishing: the design and evaluation of an embedded training email system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’07)*, 905–914. <https://doi.org/10.1145/1240624.1240760>
17. Paula M. W. Musuva, Katherine W. Getao, and Christopher K. Chepken. 2019. A new approach to modelling the effects of cognitive processing and threat detection on phishing susceptibility. *Computers in Human Behavior* 94: 154–175. <https://doi.org/10.1016/j.chb.2018.12.036>
18. Boon-Yuen Ng, Atreyi Kankanhalli, and Yunjie (Calvin) Xu. 2009. Studying users’ computer security behavior: A health belief perspective. *Decision Support Systems* 46, 4: 815–825. <https://doi.org/10.1016/j.dss.2008.11.010>
19. Gareth Norris, Alexandra Brookes, and David Dowell. 2019. The Psychology of Internet Fraud Victimisation: a Systematic Review. *Journal of Police and Criminal Psychology* 34, 3: 231–245. <https://doi.org/10.1007/s11896-019-09334-5>
20. Kathryn Parsons, Marcus Butavicius, Paul Delfabbro, and Meredith Lillie. 2019. Predicting susceptibility to social influence in phishing emails. *International Journal of Human-Computer Studies* 128: 17–26. <https://doi.org/10.1016/j.ijhcs.2019.02.007>
21. Kathryn Parsons, Marcus Butavicius, Malcolm Pattinson, Dragana Calic, Agata McCormac, and Cate Jerram. 2016. Do Users Focus on the Correct Cues to Differentiate Between Phishing and Genuine Emails? <https://doi.org/10.48550/arXiv.1605.04717>

22. Kathryn Parsons, Agata McCormac, Marcus Butavicius, and Lael Ferguson. 2010. *Human Factors and Information Security: Individual, Culture and Security Environment*. DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION EDINBURGH (AUSTRALIA) COMMAND CONTROL COMMUNICATIONS AND INTELLIGENCE DIV. Retrieved May 21, 2021 from <https://apps.dtic.mil/sti/citations/ADA535944>
23. Malcolm R. Pattinson, Cate Jerram, Kathryn Parsons, Agata McCormac, and Marcus A. Butavicius. 2011. *Managing Phishing Emails: A Scenario-Based Experiment*.
24. Jonathan Peirce, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. 2019. PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods* 51, 1: 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
25. Zoë A. Purcell, Colin A. Wastell, and Naomi Sweller. 2021. Domain-specific experience and dual-process thinking. *Thinking & Reasoning* 27, 2: 239–267. <https://doi.org/10.1080/13546783.2020.1793813>
26. Katherine A. Rawson. 2004. Exploring automaticity in text processing: Syntactic ambiguity as a test case. *Cognitive psychology* 49, 4: 333–369. <https://doi.org/10.1016/j.cogpsych.2004.04.001>
27. Nelson Siu, Lee Iverson, and Anthony Tang. 2006. Going with the flow: email awareness and task management. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work (CSCW '06)*, 441–450. <https://doi.org/10.1145/1180875.1180942>
28. J Stroop. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General* 18, 6: 643–662.
29. Valerie A. Thompson, Jamie A. Prowse Turner, Gordon Pennycook, Linden J. Ball, Hannah Brack, Yael Ophir, and Rakefet Ackerman. 2013. The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. *Cognition* 128, 2: 237–251. <https://doi.org/10.1016/j.cognition.2012.09.012>
30. Zheng Yan and Hamide Y. Gozu. 2012. Online Decision-Making in Receiving Spam Emails Among College Students. *International Journal of Cyber Behavior, Psychology and Learning (IJCBL)* 2, 1: 1–12. <https://doi.org/10.4018/ijcbpl.2012010101>