Evaluating the cognitive mechanisms of phishing detection with PEST, an ecologically valid lab-based measure of phishing susceptibility

Ziad M. Hakim¹, Natalie C. Ebner^{2,3,4,5}, Daniela S. Oliveira⁶, Sarah J. Getz^{5,7}, Bonnie E. Levin^{5,7}, Tian Lin², Kaitlin Lloyd¹, Vicky T. Lai^{1,8}, Matthew D. Grilli^{1,5}, Robert C. Wilson^{1,5,8,*}

¹ Department of Psychology, University of Arizona

² Department of Psychology, University of Florida

³ Department of Aging and Geriatric Research, Institute on Aging, University of Florida

⁴ Florida Institute for Cybersecurity, University of Florida

⁵ Evelyn F. McKnight Brain Institute

⁶ Department of Electrical and Computer Engineering, University of Florida

⁷ Department of Neurology, Miller School of Medicine, University of Miami

⁸ Cognitive Science Program, University of Arizona

* to whom correspondence should be addressed: bob@arizona.edu

Word count: 4,009

Abstract

Phishing emails constitute a major public health problem, linked to negative health outcomes due to fraud and exploitation. Because of their sheer volume and because phishing emails are designed to deceive, purely technological solutions such as filters only go so far, leaving human judgment as the last line of defense against phishing. However, in part because it is difficult to phish people under controlled laboratory conditions, little is known about the cognitive and neural mechanisms underlying phishing susceptibility. There is therefore a critical need to develop an ecologically valid measure of phishing susceptibility that can be used in the lab to test cognitive models of phishing detection. In this work, we present such a task: PEST, the Phishing Email Suspicion Test, in which participants rate a series of phishing and non-phishing emails according to their level of suspicion. By comparing suspicion scores for each email to its real-world efficacy (assessed in a field experiment in an independent group of participants), we find support for the ecological validity of PEST in that phishing emails that were more effective in the real world were more effective at deceiving people in the lab. By modeling behavior in PEST, we find evidence that the suspicion level of emails is assessed using a comparison process in which the current email is compared with previously encountered emails to determine its suspicion level. Together our task and model provide a framework for studying the cognitive neuroscience of email phishing detection in the lab.

Introduction

Email phishing is a type of cyber social engineering attack in which seemingly legitimate emails attempt to lure the receiver into performing an action with negative consequences (e.g., opening a malicious attachment that installs malware on the victims' device). While the popular conception of phishing is a message from the infamous "Nigerian Prince," modern phishing emails can be hard to distinguish from real emails, with large-scale studies suggesting click rates as high as 20% for the most effective phishing emails (PhishMe, 2016; Williams, Hinds, & Joinson, 2018). In part because of this high click rate, phishing is estimated to cost tens of billions of dollars each year (Smart People Easier, 2014) and is now also recognized as a major public health problem associated with negative health outcomes that include depression and suicide (Button, Lewis, & Tapley, 2014; Fraud Advisory Panel, 2015).

While technological solutions, such as filters and blacklists ("Google Safe Browsing," n.d.), massively reduce the number of phishing emails reaching people's inboxes, purely technical defense solutions can never be perfect. This is in part because of the sheer volume of phishing attempts, which are estimated to account for 1 in every 392 emails ("Overview of fraud and computer misuse statistics for England and Wales - Office for National Statistics," n.d.), and because of the sophistication of many phishing emails, which are carefully crafted to avoid filters. Further, phishing messages and corresponding malicious landing pages change constantly (Google, n.d.), adding challenges to machine learning methods that try to filter unwanted messages. Indeed there are even websites (such as ("Test and Optimize your Emails for the Inbox," n.d.)) where attackers can test whether a particular email will be flagged or not by several leading email providers. Thus, human decision-making is the last line of defense against phishing and there is much interest in understanding the cognitive and neural processes underlying phishing detection in order to identify those individuals who are most susceptible (Ebner et al., 2018; Lin et al., 2019; Oliveira et al., 2017) as well as to evaluate the outcome of training programs designed to reduce phishing victimization (Norris, Brookes, & Dowell, 2019).

Perhaps the Gold Standard for measuring an individual's susceptibility to phishing deception is to actually try to phish them in the home or office setting - that is, to send them a phishing email and measure whether they engage with it or not (Caputo, Pfleeger, Freeman, & Eric Johnson, 2014; Kumaraguru et al., 2009; Luo, Zhang, Burd, & Seazzu, 2013; Oliveira et al., 2017; Vishwanath, 2015; Williams et al., 2018). This direct approach has high ecological validity because the only difference from a real phishing attack is that the participants are not harmed. On the flip side, however, because the experiment occurs outside of the lab, these "field" experiments are less controlled than a lab study and any "in-the-moment" correlates of phishing susceptibility (such as physiological and neural measures) are almost impossible to obtain. Another challenge is that the number of emails sent to participants in field experiments is often small because of the need to simulate real-world phishing, where emails might be spaced out over several days to avoid arousing suspicion. This results in limited sampling of each person's susceptibility to attack, and this small amount of data per person makes it difficult to test cognitive models of how individual phishing emails are evaluated and detected.

Another approach is to try to measure phishing susceptibility in the lab. In the most direct version of this approach, the experimenter tries to phish people in the lab, for example by having them browse real and phishing websites to see whether or not they divulge sensitive information (Gavett et al., 2017). Other lab-based experiments involve roleplaying another person checking

their emails (Downs, Holbrook, & Cranor, 2007; Sheng, Holbrook, Kumaraguru, Cranor, & Downs, 2010) or rating a series of emails according to how suspicious they appear or how likely the person would be to respond to such an email (Alsharnouby, Alaca, & Chiasson, 2015; Dhamija, Tygar, & Hearst, 2006; Jones, Towse, Race, & Harrison, 2019; Kelley & Bertenthal, 2016; Rajivan & Gonzalez, 2018; Wood, Liu, Hanoch, Xi, & Klapatch, 2018; Yan & Gozu, 2012). Despite the increased experimental control offered by these lab-based tasks, the extent to which these measures are ecologically valid remains unknown. Given the urgency and magnitude of this growing crisis from phishing on health and well-being, there is a need to document and quantify the relationship between lab and real-world behavior.

In this paper we provide initial evidence that lab and real-world phishing susceptibility are related. Using a new task, the Phishing Email Suspicion Test (PEST), we asked participants to rate the degree to which a series of 160 emails appeared suspicious. By design, 84 of these emails had been previously used in a field-experimental phishing study using an independent sample of participants (Oliveira et al., 2017). This allowed us to determine whether emails that were effective in real life (i.e., people had fallen for them in the field experiment), were also effective in the lab (i.e., were rated as *less* suspicious in PEST).

In addition to testing the ecological validity of phishing detection in the lab, PEST allowed us to probe the cognitive mechanisms underlying phishing detection using a computational model. This model was inspired by the similarity between PEST, in which participants evaluate a series of emails according to their level of suspicion, and more traditional tasks from psychophysics, in which participants evaluate a series of physical stimuli (e.g., loudness of sounds or heaviness of weights) according to their perceived properties. Such psychophysical tasks have been extensively studied for over a hundred years, and there are well-established models as to how people make the transformation from stimulus to rating.

One model in particular (Ward & Lockhead, 1970) suggests that people assign a rating to stimuli by comparing it with previous stimuli. For example, in the case of a loudness rating task, participants compare the loudness of the current stimulus to that of the previous stimulus. If they judge the current stimulus to be louder, they give it a higher rating; if they judge the current stimulus to be quieter, they give it a lower rating. Such a comparison process reveals itself in the form of sequential effects in the ratings, whereby the current rating is biased by an “assimilative” effect of the past rating (i.e., the current rating is biased towards the past rating), and a “contrastive” effect of the past stimulus (i.e., the current rating is biased away from the past stimuli) (Ward & Lockhead, 1970).

We reasoned that email evaluation may involve a similar comparison process, such that people judge how suspicious an email is by comparing it to previous emails. To test this idea, we used a linear regression model to quantify sequential effects in PEST. In particular, we asked whether there was a positive (assimilative) effect of past rating and a negative (contrastive) effect of past stimulus. More generally, this formal modeling approach allowed us to quantify individual differences in phishing detection in terms of the regression parameters, a potentially more robust measure than “accuracy” because it does not depend on the exact emails participants had seen during the lab task.

Methods

Participants

Ninety seven undergraduates (36 male, 61 female, ages 18-27, mean 18.9, standard deviation 1.39) were recruited from the University of Arizona Psychology subject pool. An additional three participants completed the task as part of their enrollment in the subject pool but were excluded from the analysis because they were under 18 (as required by our subject pool-IRB agreement). The large sample size was designed to give us high statistical power while still being feasible to run within one semester. Participants were not paid for participation, but received course credit for completing the experiment. The study was approved by the Institutional Review Board at the University of Arizona and all participants provided written informed consent.

The Phishing Email Suspicion Test (PEST)

Participants were presented with a series of 160 emails and instructed to classify on a four-point scale from 1, “definitely safe,” to 4, “definitely suspicious” via keyboard press (see Figure 1 for schematic display and sample emails). Participants were told to maximize their classification accuracy in order to maximize their “score,” although they received no monetary or other reward for their performance. The order of emails seen in the task was randomized for each participant. To minimize learning during the task, participants did not receive feedback about their performance until the end of the experiment. The task was coded using Psychtoolbox for Matlab (Brainard, 1997). Key press as well as response times were recorded.

Emails varied in their source (real emails vs simulated emails created by us) and legitimacy (safe vs phishing) such that there were four types of email (real-safe, simulated-safe, real-phish, and simulated-phish). Each participant saw 40 examples of each type of email, but because we had more than 40 examples of some emails types, they did not all see the same emails. Overall the 160 emails seen by participants were sampled from 348 emails made up of 140 real phishing emails, 84 simulated phishing emails, 84 simulated real emails and 40 real emails.

Real phishing emails were sampled from a set of one hundred and forty genuine phishing emails obtained from four websites that collect phishing examples (“Information Security at UVA,” n.d., “Phishing and Scam Emails: A Realtime Database of Phishing Emails,” n.d.; Information Security University of Arizona, n.d.; UCLA Information Security Office, n.d.). These emails were edited to improve relevance for University of Arizona students only in cases where they mentioned specific institutions or locations (e.g., University of Virginia was changed to University of Arizona). The real phishing emails were selected such that they spanned a range of effectiveness, from obvious phishing to more convincing attacks (based on subjective judgments of two of the authors, ZMH and KL). Simulated phishing emails were taken from the eighty four emails used in the PHishing Internet Task (PHIT) (Lin et al., 2019; Oliveira et al., 2017), see below for details.

Safe emails also comprised real and simulated messages. Forty real safe emails were taken from our inboxes and included messages from banks and paypal, as well as password resets or general account management. Eighty-four simulated safe emails were adapted versions of the original PHIT emails that were altered to make them seem less suspicious, but matched on other features such as word count, purported identity of the sender and topic. In particular we corrected for poor syntax, grammatical errors, informal language, and aggressive tone in the original PHIT emails.

Google	
Gmail	Problems with item delivery, n.01403407
Compose Inbox (160) Starred Important Drafts More ...	<p>Dear Customer,</p> <p>We can not deliver your parcel arrived at August 17.</p> <p>Please review delivery label here: www.ups.com/track/label-29382</p> <p>Best regards,</p> <p>Sergio James,</p> <p>UPS Senior Support Manager</p>
	<div>Definitely Safe</div> <div>Possibly Safe</div> <div>Possibly Suspicious</div> <div>Definitely Suspicious</div>

SIMULATED PHISH	REAL SAFE	SIMULATED SAFE
Free Gift Received From Matt	Please confirm your e-mail address - Airbnb	Apple iCloud Security
Dear Bob,	Hi Bob,	Dear Bob,
Your friend Matt sent you a free gift through Games with Friends, a new social media site for keeping in touch through games and trades.	Welcome to Airbnb! In order to get started, you need to confirm your email address.	Your AppleID was used to sign in to iCloud on an iPhone 6S in Tucson, Arizona at 7:06 PM on Monday, January 8, 2018.
Check out what Matt got you by following the link below:	Thanks,	If this was your sign-in, please disregard this email. If this was not your sign in, please contact Apple Support at the link below.
http://www.tucsonweekend.com/gift-social/	Confirm Email	https://support.apple.com
Matt Games with Friends	The Airbnb Team Sent with Love from Airbnb Airbnb, Inc., 888 Brannan St, San Francisco, CA 94103	iCloud Authentication Services Copyright 2018 Apple Inc. All Rights Reserved

Figure 1 - The Phishing Email Suspicion Test (PEST). Participants were presented with a series of 160 emails from one of four categories (real-safe, simulated-safe, real-phish, simulated-phish, 40 emails per category) in randomized order and instructed to rate them on a four-point scale from “definitely safe” to “definitely suspicious.” (Top) Schematic of the display seen by participants when evaluating an email. In this case the email is a real phishing email. (Bottom) Examples of the subject line and text from the other types of email: a simulated phishing email, a real safe email, and a simulated safe email.

Real-world phishing efficacy for a subset of emails

Phishing efficacy of the 84 simulated phishing emails had previously been assessed in a field experiment by our group (Lin et al., 2019; Oliveira et al., 2017). In this Phishing Internet Task (PHIT), 158 participants were sent emails to the email address they had registered with the study, and a browser plug in recorded whether participants clicked on the link embedded in each email. This data allowed us to construct a measure of phishing efficacy for each email had it occurred in real life, as the proportion of times an email was successful in eliciting a click. In particular, each participant received 21 emails each during the course of 21 days (one per day). Emails had been created based on a large set of authentic spam emails from an independent sample of Internet users (Oliveira et al., 2019). On average each email was sent to just under 40 participants (range 37-42).

Results

Participants were more suspicious of phishing than safe emails

Overall, participants performed above chance in PEST, achieving a mean accuracy (defined as the safe and phishing emails were correctly classified as possibly or definitely safe or phishing) of 62%. Regarding mean suspicion scores for each of the four categories of emails (real-safe, simulated-safe, real-phish, simulated-phish), we conducted a repeated measures ANOVA with email legitimacy (safe/phish) and source (real/simulated) as within-subject factors (Supplementary Table 1). This analysis revealed a main effect of legitimacy ($F(1, 288) = 455.1$, $p < 0.001$, partial $\eta^2 = 0.61$), a main effect of source ($F(1, 288) = 56.3$, $p < 0.001$, partial $\eta^2 = 0.16$), and an interaction effect between legitimacy and source ($F(1, 285) = 66.5$, $p < 0.001$, partial $\eta^2 = 0.19$). Additional analysis further showed that these effects were due to participants rating phishing emails as more suspicious than safe emails ($t(96) = 20.3$, $p < 0.001$), simulated-safe emails as more suspicious than real-safe emails ($t(96) = 11.3$, $p < 0.001$), and simulated-phishing emails the same as real-phishing emails ($t(96) = 0.48$, $p = 0.63$) (Figure 2A). Despite this difference in overall score between real and simulated safe emails, ratings for both types of simulated emails were highly correlated with ratings of real emails (Figure 2B, for safe emails $r(95) = 0.56$, $p < 0.001$; Figure 2C for phishing emails $r(95) = 0.45$, $p < 0.001$).

Item analysis reveals a wide range of suspicion scores across emails

Averaging across participants allowed us to compute a suspicion score for each email. As shown in Figure 3, there was a wide range of suspicion scores for all four types of email. While safe emails were generally rated less suspicious than phishing emails, there was considerable overlap in the ratings for the four types of emails. For example, the least suspicious phishing email (mean score = 1.46) was ranked almost as safe as the least suspicious safe email (mean score = 1.38). Likewise the most suspicious safe email (mean score = 3.26) was rated more suspicious than the average phishing email (mean score = 2.79). This overlap in suspicion score between phishing and safe emails illustrates how convincing some phishing emails can be.

Figure 2 - Suspicion scores in PEST. Safe emails are depicted in blue and phishing emails are depicted in red. Each dot represents the mean suspicion score given to each type of email by one participant. (A) Mean suspicion score for each of the four different email types (real-safe, simulated-safe, real-phish, simulated-phish). Phishing emails were rated as more suspicious than safe emails and simulated safe emails were rated as more suspicious than real safe emails. Notched box plots represent median, confidence intervals for the median, upper and lower quartiles of suspicion scores and range. * indicates $p < 0.001$. (B, C) Correlation between suspicion scores for real and simulated emails. Each participant corresponds to a single dot. The correlations between simulated and real safe (B) and simulated and real phishing (C) emails suggest that performance on the simulated emails was positively associated with performance on real emails, for both safe and phishing emails.

Figure 3 - Average suspicion scores for each email type. Each email corresponds to a different dot. Emails were grouped by type (real-safe, simulated-safe, real-phish, simulated phish). Safe emails are depicted in blue and phishing emails in red. While, on average, safe emails were ranked as less suspicious, there was considerable overlap in the mean suspicion score for each type of email.

More effective phishing emails were rated as less suspicious

As a test of the ecological validity of our newly developed PEST, we compared the average suspicion score of each email on PEST to the efficacy of the email as previously assessed in the original PHIT field experimental (Lin et al., 2019; Oliveira et al., 2017). By design, 84 emails (the simulated phishing emails) for which we had real-world measures of efficacy were included in PEST. As shown in Figure 4A, we found a negative correlation between PEST and real-world behavior (Spearman's $\rho(82) = -0.22$, $p = 0.048$) such that emails with a low suspicion score in the lab had been more effective at phishing participants in the real world field experiment.

As a complement to the correlation analysis, we ran a separate analysis in which we split the emails into two groups depending on whether they had been clicked on at all (“clicked”, $n = 44$ emails) or not (“not-clicked”, $n = 40$ emails) in PHIT. Consistent with the correlation analysis, clicked emails were rated as less suspicious than not-clicked emails ($t(82) = 2.09$, $p = 0.04$) (Figure 4B). Together these findings provide initial evidence that PEST may be an ecologically valid measure of phishing behavior in the real world.

Figure 4 - Ecological validity of PEST. (A) Correlation between PEST behavior and real-world efficacy of each email in the PHIT field experiment, operationalized as the fraction of times the link in the email was clicked. Emails rated as more suspicious in PEST had been more effective at phishing people in the real world. (B) Alternate analysis separating emails into two groups based on whether they had been successful at phishing at least one Internet user in the PHIT field experiment (i.e., at least one person clicked on the link embedded in the email; “clicked” email) or not (no person had clicked on the link embedded in the email; “not-clicked” email). In line with the correlation analysis, emails that were not clicked in the original PHIT were also rated as more suspicious in PEST. * indicates $p < 0.05$.

Regression model of PEST behavior

To further analyze PEST behavior, we built a linear regression model similar to that used in (Wilson, 2018)). This model assumes that participants' ratings for each email are computed according to

$$c_t = \beta_{bias} + \beta_f f_t + \beta_{fPast} f_{t-1} + \beta_{cPast} c_{t-1} \quad (\text{Equation 1})$$

where c_t is the rating on trial t , f_t is the current stimulus “strength,” computed as the average suspicion score across participants for this stimulus, f_{t-1} is the previous stimulus, and c_{t-1} is the previous choice. The regression weights are the free parameters of the model and denote the overall phish bias (β_{bias}), the effect of the current stimulus (β_f), the effect of the last stimulus (β_{fPast}), and the effect of the last rating (β_{cPast}). For the purposes of the regression (and to keep the notation in line with our previous work; (Wilson, 2018)), the ratings and suspicion scores were normalized to range between -1 and +1.

Regression weights were estimated separately for each participant using a standard least-squares method (implemented via the `glmfit` function in Matlab). Results are shown in Figure 5. The largest effect was for the current stimulus with a median regression weight close to 1, consistent with participants' score being correlated with the average score of other people, but with considerable variability across people, consistent with individual differences in phishing susceptibility. This effect of the current stimulus was positive, assimilative, for all participants and statistically significant ($t(96) = 27.28$, $p < 0.001$). There was also a significant phish bias. That is, all else being equal, participants were more likely to say that an email was suspicious than safe ($t(96) = 4.34$, $p < 0.001$). More subtly, we also saw evidence for sequential effects in the form of a negative, contrastive, effect of the last stimulus ($t(96) = -3.53$, $p < 0.001$) and a positive, assimilative, effect of the last rating ($t(96) = 4.40$, $p < 0.001$). These effects are consistent with previous work in sequential judgment tasks with no feedback and suggest a comparison process for email evaluation in which the current email is evaluated relative to the last (Holland & Lockhead, 1968).

Figure 5 - Regression model of PEST behavior. Regression weights for the overall phish bias (β_{bias}), the effect of the current stimulus (β_f), the effect of the past stimulus (β_{fPast}), and the effect of the past rating (β_{cPast}) are plotted for each participant. Ratings are most strongly influenced by the average suspicion score of the current email (blue), but also show an overall phish bias (red). In addition we see evidence of sequential effects in the form of a contrastive effect of past stimulus (yellow) and an assimilative effect of past rating (green). * indicates $p < 0.001$.

Discussion

In this paper, we introduced PEST as a tool for assessing email phishing susceptibility in the lab with ecological validity. In this task participants evaluated a series of 160 emails regarding their level of suspicion. While participants were generally more suspicious of phishing than safe emails, their performance was far from perfect, with participants only classifying emails correctly about 62% of the time. Such a relatively low accuracy is consistent with previous reports, such as (Jones et al., 2019) who found an accuracy of 68% in a similar task. This low accuracy is indicative of how convincing phishing emails can be. Indeed, when we looked at the average suspicion score given to each email (Figure 3), there was considerable overlap between the suspicion score given to safe vs. phishing emails, with one of the real phishing emails being among the least suspicious of all.

A unique feature of PEST is that the real-world efficacy of 84 simulated-phishing emails had been previously assessed in a field experiment (Lin et al., 2019; Oliveira et al., 2017). Thus, by comparing real-world and lab-based efficacy, we were able to test, for the first time, whether email phishing susceptibility in the lab is related to email phishing susceptibility in the real world. In particular, we found that emails that were more effective at phishing people in the real world were rated as less suspicious in the lab, suggesting that PEST can capture at least some of the processes underlying phishing susceptibility in real life. This finding of possible ecological validity, in turn, lends weight to the idea of using PEST in combination with neuroimaging and other physiological measures (e.g., eye tracking and galvanic skin response) in the lab, to capture the cognitive and neural processes of real-life phishing decision making.

One limitation of our current approach is that the lab-based and real-world measures of email phishing susceptibility were acquired from different participants. This limits the scope of our conclusions such that, while we can say that PEST is effective in assessing the efficacy of individual *emails*, we cannot say that PEST is effective at determining the phishing susceptibility of individual *people*. To determine this, we will need a within-subjects design to examine PEST in parallel with a direct phishing measure. These data would allow us to assess the extent to which PEST behavior captures real-world phishing behavior, and which components of PEST (e.g., regression coefficients in Equation 1) correlate with real-world phishing decision-making across participants.

In addition to assessing the ecological validity of PEST, the large number of trials in our experiment allowed us to model trial-by-trial behavior with linear regression. We built a model in which participant ratings were determined by a combination of the current email, a bias towards saying an email is suspicious, and sequential effects from the last email and its rating. Of particular interest were the two sequential effects: a contrastive effect for past stimulus and an assimilative effect for past rating. While sequential effects have not previously been reported in phishing tasks, they are known to occur in other sequential judgment tasks without performance feedback (Garner, 1953; Holland & Lockhead, 1968; Parducci & Marshall, 1962; Pegors, Mattar, Bryan, & Epstein, 2015; Ward & Lockhead, 1970; Wedell, Parducci, & Edward Geiselman, 1987), see (Kiyonaga, Scimeca, Bliss, & Whitney, 2017) for review.

Guided by this previous literature, the presence of these sequential effects in PEST suggests a model for real-world phishing decision making. In particular, one way in which contrastive and assimilative sequential effects can simultaneously occur is if the judgment is made via a comparison process with recent stimuli (Holland & Lockhead, 1968). In the simplest form of this

model, which considers only the stimulus and rating on the last trial, the rating on the present trial is computed by adjusting the rating from the last trial in proportion to the difference between the present stimulus and the last stimulus. Written mathematically this implies that

$$c_t = c_{t-1} + s_t + s_{t-1} \quad (\text{Equation 2})$$

hence the past rating, c_{t-1} , has a positive, i.e. assimilative, effect on the current rating, while the past stimulus, s_{t-1} , has a negative, i.e. contrastive, effect. More generally, if the comparison process is not just with one but the average of several past stimuli, then the rating takes the form

$$c_t = s_t + \sum_i w_i (c_{t-i} - s_{t-i}) \quad (\text{Equation 3})$$

where w_i describes the weighting given to the stimulus at time $t - i$.

In line with this equation, including more terms in the regression analysis suggests that the assimilative effect of past rating and the contrastive effect of past stimuli extend multiple trials into the past (Supplementary Figure 2). Could real-world phishing decisions made in a similar manner? That is, are phishing decisions made by comparing the current email with exemplars of safe and phishing emails we have seen before? Such a process is certainly possible, but would need more extensive field experiments to test, for example by sending emails in pairs such that the effect of neighboring emails in the inbox on susceptibility to a phishing email could be assessed.

Clearly more work is needed to understand the complex set of processes that underlie one's decision to engage with a phishing email or not. This study has taken initial steps toward the development of a tool for in-lab assessment of email phishing susceptibility, which is efficient in administration and allows data collection to inform the cognitive mechanisms of phishing detection.

Acknowledgements

This work was supported by a pilot grant from the McKnight Brain Research Foundation, NSF grant SBE-1450624, and NIH grant 1R01AG057764-01A1.

Author contributions

RCW, MDG, NCE, DSO, SJG, and BEL conceived the study. ZMH built the experiment with supervision from RCW and MDG. ZMH and KL designed the simulated-safe emails. NCE, DSO, and TL provided the original simulated phishing emails and the PHIT data. ZMH ran data collection with supervision from RCW and MDG. ZMH and RCW analyzed the data. RCW wrote the paper with input from MDG, NCE, DSO, SJG, BEL, ZMH, and VL.

References

- Alsharnouby, M., Alaca, F., & Chiasson, S. (2015). Why phishing still works: User strategies for combating phishing attacks. *International Journal of Human-Computer Studies*, Vol. 82, pp. 69–82. <https://doi.org/10.1016/j.ijhcs.2015.05.005>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, Vol. 10, pp. 433–436. <https://doi.org/10.1163/156856897x00357>
- Button, M., Lewis, C., & Tapley, J. (2014). Not a victimless crime: The impact of fraud on individual victims and their families. *Security Journal*, Vol. 27, pp. 36–54. <https://doi.org/10.1057/sj.2012.11>

- Caputo, D. D., Pfleeger, S. L., Freeman, J. D., & Eric Johnson, M. (2014). Going Spear Phishing: Exploring Embedded Training and Awareness. *IEEE Security & Privacy*, Vol. 12, pp. 28–38. <https://doi.org/10.1109/msp.2013.106>
- Deasley, S., Evans, M. A., Nowak, S., & Willoughby, D. (2018). Sex Differences in Emergent Literacy and Reading Behaviour in Junior Kindergarten. *Canadian Journal of School Psychology*, Vol. 33, pp. 26–43. <https://doi.org/10.1177/0829573516645773>
- Dhamija, R., Tygar, J. D., & Hearst, M. (2006). Why phishing works. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '06*. <https://doi.org/10.1145/1124772.1124861>
- Downs, J. S., Holbrook, M., & Cranor, L. F. (2007). Behavioral response to phishing risk. *Proceedings of the Anti-Phishing Working Groups 2nd Annual eCrime Researchers Summit on - eCrime '07*. <https://doi.org/10.1145/1299015.1299019>
- Ebner, N. C., Ellis, D. M., Lin, T., Rocha, H. A., Yang, H., Dommaraju, S., ... Oliveira, D. S. (2018). Uncovering Susceptibility Risk to Online Deception in Aging. *The Journals of Gerontology. Series B, Psychological Sciences and Social Sciences*. <https://doi.org/10.1093/geronb/gby036>
- Fraud Advisory Panel. (2015). Supporting the victims of fraud: The year in review 2014–2015. Retrieved August 24, 2019, from <https://www.fraudadvisorypanel.org/wp-content/uploads/2015/07/FAP-Yearly-Review-2015-WEB.pdf>
- Garner, W. R. (1953). An informational analysis of absolute judgments of loudness. *Journal of Experimental Psychology*, 46(5), 373–380.
- Gavett, B. E., Zhao, R., John, S. E., Bussell, C. A., Roberts, J. R., & Yue, C. (2017). Phishing suspiciousness in older and younger adults: The role of executive functioning. *PloS One*, 12(2), e0171620.
- Google. (n.d.). Disclosing vulnerabilities to protect users across platforms. Retrieved October 10, 2019, from Google Online Security Blog website: <https://security.googleblog.com/2019/03/disclosing-vulnerabilities-to-protect.html>
- Google Safe Browsing. (n.d.). Retrieved October 10, 2019, from <https://safebrowsing.google.com/>
- Holland, M. K., & Lockhead, G. R. (1968). Sequential effects in absolute judgments of loudness. *Perception & Psychophysics*, Vol. 3, pp. 409–414. <https://doi.org/10.3758/bf03205747>
- Information Security at UVA. (n.d.). Retrieved August 24, 2019, from Information Security Alerts & Warnings website: <https://security.virginia.edu/security-alerts-and-warnings>
- Information Security University of Arizona. (n.d.). Phishing Alerts. Retrieved August 24, 2019, from https://security.arizona.edu/phishing_alerts
- Jones, H. S., Towse, J. N., Race, N., & Harrison, T. (2019). Email fraud: The search for psychological predictors of susceptibility. *PloS One*, 14(1), e0209684.
- Kelley, T., & Bertenthal, B. I. (2016). Attention and past behavior, not security knowledge, modulate users' decisions to login to insecure websites. *Information and Computer Security*, Vol. 24, pp. 164–176. <https://doi.org/10.1108/ics-01-2016-0002>
- Kiyonaga, A., Scimeca, J. M., Bliss, D. P., & Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences*, 21(7), 493–497.
- Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M. A., & Pham, T. (2009). School of phish. *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09*. <https://doi.org/10.1145/1572532.1572536>

- Lin, T., Capecci, D. E., Ellis, D. M., Rocha, H. A., Dommaraju, S., Oliveira, D. S., & Ebner, N. C. (2019). Susceptibility to Spear-Phishing Emails. *ACM Transactions on Computer-Human Interaction*, Vol. 26, pp. 1–28. <https://doi.org/10.1145/3336141>
- Luo, X. (robert), Zhang, W., Burd, S., & Seazzu, A. (2013). Investigating phishing victimization with the Heuristic–Systematic Model: A theoretical framework and an exploration. *Computers & Security*, Vol. 38, pp. 28–38. <https://doi.org/10.1016/j.cose.2012.12.003>
- McGeown, S., Goodwin, H., Henderson, N., & Wright, P. (2012). Gender differences in reading motivation: does sex or gender identity provide a better account? *Journal of Research in Reading*, Vol. 35, pp. 328–336. <https://doi.org/10.1111/j.1467-9817.2010.01481.x>
- Norris, G., Brookes, A., & Dowell, D. (2019). The Psychology of Internet Fraud Victimization: a Systematic Review. *Journal of Police and Criminal Psychology*. <https://doi.org/10.1007/s11896-019-09334-5>
- Oliveira, D., Ebner, N., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., ... Lin, T. (2017). Dissecting Spear Phishing Emails for Older vs Young Adults. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. <https://doi.org/10.1145/3025453.3025831>
- Overview of fraud and computer misuse statistics for England and Wales - Office for National Statistics. (n.d.). Retrieved October 10, 2019, from <https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/overviewof fraudandcomputermisusestatisticsforenglandandwales/2018-01-25>
- Parducci, A., & Marshall, L. M. (1962). Assimilation vs. contrast in the anchoring of perceptual judgements of weight. *Journal of Experimental Psychology*, 63, 426–437.
- Pegors, T. K., Mattar, M. G., Bryan, P. B., & Epstein, R. A. (2015). Simultaneous perceptual and response biases on sequential face attractiveness judgments. *Journal of Experimental Psychology. General*, 144(3), 664–673.
- Phishing and Scam Emails: A Realtime Database of Phishing Emails. (n.d.). Retrieved August 24, 2019, from Phishing and Scam Emails: A Realtime Database of Phishing Emails website: <https://philmir.wordpress.com>
- PhishMe. (2016). Enterprise Phishing Susceptibility Report. Retrieved May 30, 2019, from https://cofense.com/wp-content/uploads/2017/10/PhishMe_EnterprisePhishingSusceptibilityReport_2015_Final.pdf
- Quinn, J. M., & Wagner, R. K. (2015). Gender Differences in Reading Impairment and in the Identification of Impaired Readers: Results From a Large-Scale Study of At-Risk Readers. *Journal of Learning Disabilities*, 48(4), 433–445.
- Rajivan, P., & Gonzalez, C. (2018). Creative Persuasion: A Study on Adversarial Behaviors and Strategies in Phishing Attacks. *Frontiers in Psychology*, Vol. 9. <https://doi.org/10.3389/fpsyg.2018.00135>
- Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L. F., & Downs, J. (2010). Who falls for phish? *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10*. <https://doi.org/10.1145/1753326.1753383>
- Smart People Easier. (2014). Smart people easier to scam. Retrieved August 24, 2019, from http://http://www.ultrascan-agi.com/public_html/html/pdf_files/Pre-Release-419_Advance_Fee_Fraud_Statistics_2013-July-10-2014-NOT-FINAL-1.pdf
- Test and Optimize your Emails for the Inbox. (n.d.). Retrieved August 24, 2019, from <https://glockapps.com/spam-testing/>

- UCLA Information Security Office. (n.d.). PHISH BOWL/PHISHING SCAMS. Retrieved August 24, 2019, from <https://www.it.ucla.edu/security/alerts/phishing-scams>
- van den Brink, D., Van Berkum, J. J. A., Bastiaansen, M. C. M., Tesink, C. M. J. Y., Kos, M., Buitelaar, J. K., & Hagoort, P. (2012). Empathy matters: ERP evidence for inter-individual differences in social language processing. *Social Cognitive and Affective Neuroscience*, 7(2), 173–183.
- 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*, Vol. 20, pp. 570–584. <https://doi.org/10.1111/jcc4.12126>
- Ward, L. M., & Lockhead, G. R. (1970). Sequential effects and memory in category judgments. *Journal of Experimental Psychology*, Vol. 84, pp. 27–34. <https://doi.org/10.1037/h0028949>
- Wedell, D. H., Parducci, A., & Edward Geiselman, R. (1987). A formal analysis of ratings of physical attractiveness: Successive contrast and simultaneous assimilation. *Journal of Experimental Social Psychology*, Vol. 23, pp. 230–249. [https://doi.org/10.1016/0022-1031\(87\)90034-5](https://doi.org/10.1016/0022-1031(87)90034-5)
- Williams, E. J., Hinds, J., & Joinson, A. N. (2018). Exploring susceptibility to phishing in the workplace. *International Journal of Human-Computer Studies*, Vol. 120, pp. 1–13. <https://doi.org/10.1016/j.ijhcs.2018.06.004>
- Wilson, R. C. (2018). *Sequential choice effects predict prevalence-induced concept change*. <https://doi.org/10.31234/osf.io/75bpy>
- Wood, S., Liu, P.-J., Hanoch, Y., Xi, P. M., & Klapatch, L. (2018). Call to claim your prize: Perceived benefits and risk drive intention to comply in a mass marketing scam. *Journal of Experimental Psychology: Applied*, Vol. 24, pp. 196–206. <https://doi.org/10.1037/xap0000167>
- Yan, Z., & Gozu, H. Y. (2012). Online Decision-Making in Receiving Spam Emails Among College Students. *International Journal of Cyber Behavior, Psychology and Learning*, Vol. 2, pp. 1–12. <https://doi.org/10.4018/ijcbpl.2012010101>