Unrealistic Promises and Urgent Wording Differently Affect Suspicion of Phishing and Legitimate Emails

McKenna K. Tornblad[1], Miriam E. Armstrong[1], Keith S. Jones[1], and Akbar Siami Namin[2]
[1]Dept. of Psychological Sciences, [2]Dept. of Computer Science, Texas Tech University, Lubbock, TX

Phishing emails have certain characteristics, including wording related to urgency and unrealistic promises (i.e., "too good to be true"), that attempt to lure victims. To test whether these characteristics affected users' suspiciousness of emails, users participated in a phishing judgment task in which we manipulated 1) email type (legitimate, phishing), 2) consequence amount (small, medium, large), 3) consequence type (gain, loss), and 4) urgency (present, absent). We predicted users would be most suspicious of phishing emails that were urgent and offered large gains. Results supporting the hypotheses indicate that users were more suspicious of phishing emails with a gain consequence type or large consequence amount. However, urgency was not a significant predictor of suspiciousness for phishing emails, but was for legitimate emails. These results have important cybersecurity-related implications for penetration testing and user training.

Phishing attacks are malicious electronic messages distributed to many people with the intent to gain access to personal information (Butavicius et al., 2015). Phishing attacks are increasing in prevalence, success, and impact across the globe. These attacks are also highly concerning for IT practitioners (CyberEdge Group, 2020) as they are increasingly costing companies and individuals through data loss, account compromise, malware infections, and direct financial loss (APWG, 2021; Proofpoint, 2021).

Phishing Email Characteristics

Previous studies investigating phishing have found that phishing emails tend to have certain characteristics (Downs et al., 2006; Lötter & Futcher, 2015; Wang et al., 2012), including 1) urgent wording and 2) unrealistic promises. Urgent wording could include ultimatums, such as "respond immediately." Unrealistic promises could include offers that are "too good to be true," such as winning a large sum of money from a lottery. An example of these two features combined could be "respond in the next 24 hours to claim your $5,000 lottery winnings."

Attackers use these features to motivate users to respond and attack success may be due to social psychological influences. Specifically, falling for a phishing scam may be due to the lack of users' cognitive involvement and information processing that occurs when users choose heuristic evaluation methods rather than deliberative, conscious evaluations. If cues are not attended to or elaborated on, users are less likely to recognize them and thus, may be more likely to fall for a phishing attack (Cialdini, 2001; Vishwanath et al., 2011). By including language that evokes emotions such as fear ("respond immediately") or excitement ("claim your $5,000"), scammers attempt to circumvent users' conscious evaluation of phishing cues and increase the likelihood that they will fall for the scam.

Although attackers may attempt to distract users from noticing these phishing cues, both experts and novices are aware that phishing emails often include urgent wording and unrealistic promises (Zielinska et al., 2015). Others have argued that continuing to use phishing tactics with which users are familiar could lead to these attacks being less effective (Downs et al., 2006). Therefore, if an email urges a user to respond quickly for a large reward, this may lead them to suspect that they are being phished.

Current Study

To investigate that possibility, we had participants judge whether 24 emails were phishing emails. Half (12) were legitimate emails and half (12) were phishing emails. Both sets varied in terms of 1) whether they mentioned a gain or a loss as a consequence, 2) the amount of the consequence, and 3) whether it was urgent or not.

We predicted that the presence of unrealistic promises (i.e., a large gain) and urgent wording would cause users to suspect that an email was phishing. Therefore, we predicted that users would be more likely to judge an email as phishing if it had the characteristics of *gain* (out of consequence type), *large* (out of consequence amount), and *urgent* (out of urgency). Although we predicted this for phishing emails, it could also be true for legitimate emails, so we also investigated that possibility.

The current study has important cybersecurity implications if our predictions are found to be correct. For example, penetration testing campaigns that include only phishing emails with urgency and unrealistic promises may only be identifying the most susceptible users. Although other users may be able to pick up on these often-used cues and avoid falling for the phish, they may be vulnerable to other kinds of phishing emails. Further, users should be aware, and anti-phishing training programs should emphasize, that there may be a tendency to focus on urgency and unrealistic promises when evaluating potential phishing emails, which may unwittingly cause users to fall for phishing emails that do not utilize those cues.

METHOD

Study Design

This study employed a 2 (email type: legitimate, phishing) x 2 (consequence type: gain, loss) x 3 (consequence amount: small, medium, large) x 2 (urgency: present, absent) full factorial within-subjects design. Materials and procedures were approved by an Internal Review Board.

Participants

Participants were 767 students given course credit for their participation. This initial data set was cleaned to remove individuals who did not complete the study, self-reported that their data should not be used, or spent under 20 or over 65 minutes on the study. Based on pilot testing, we concluded that participants who completed the study in under 20 minutes or over 65 minutes likely responded carelessly or were not devoting their full attention to the study, respectively. Participants were also excluded if they were missing phishing judgments for any of the 24 emails.

The resulting dataset consisted of 353 participants, 265 (75.1%) of which were female. Additionally, one participant did not identify as male or female. Their gender was treated as missing to include them in all other analyses. Participants' ages ranged from 16 to 49 years (M = 20.04, SD = 3.47). Participants reported being fluent in English and frequently using email.

Materials

Twenty-four email stimuli were created for this experiment: 12 legitimate emails (see Figure 1 for an example) and 12 phishing emails (see Figure 2 for an example). The general process by which these stimuli were developed was modeled after that used in previous phishing work (Downs et al., 2006).

Figure 1

Sample Legitimate Email

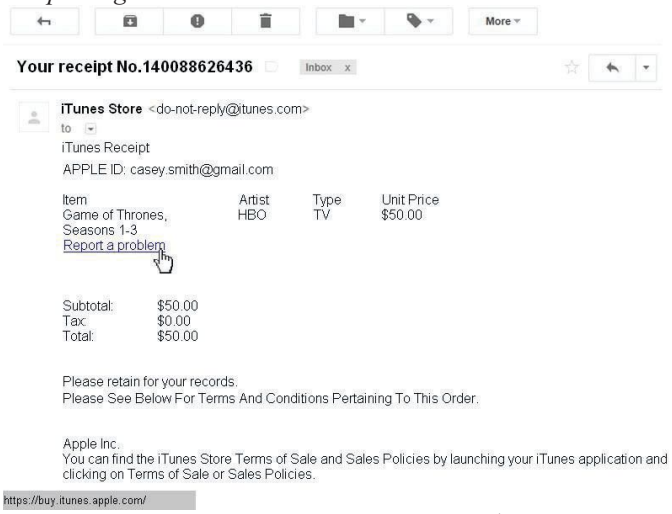

Note. This email contains a *loss* consequence, a *medium* [$50] consequence amount, and is *not urgent*.

Figure 2

Sample Phishing Email

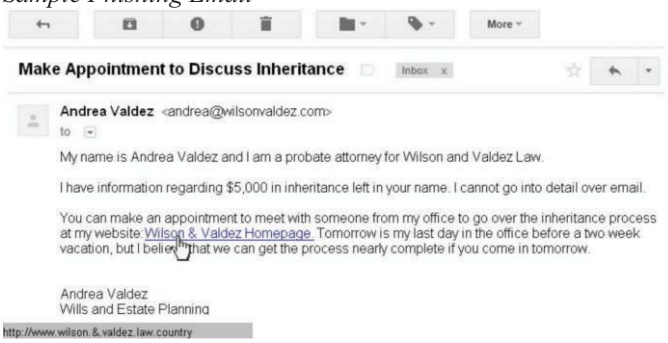

Note. This email contains a *gain* consequence, a *large* [$5,000] consequence amount, and is *urgent*.

Each of the 24 emails had certain characteristics. Specifically, each stimulus was an image of an opened email in a Gmail browser that was sent to a fictitious persona ("Casey Smith") participants adopted, consistent with previous studies (Canfield et al., 2016). Further, all body messages were 50-100 words long and contained one hyperlink; in the email image, a mouse hovered over the link so that a URL was visible at the bottom of the email (Canfield et al., 2016).

The emails' sender address, subject line, and main body, which were modeled after legitimate emails found online and in the researchers' inboxes, were modified to accommodate the experimental manipulations: email type (legitimate, phishing), consequence amount (small [$5], medium [$50], large [$5,000]), consequence type (gain, loss), and urgency (present, absent). The following paragraphs detail the differences between emails due to these four manipulations.

Emails were one of two types: legitimate or phishing emails. Legitimate emails contained some personal information related to a fictional email recipient, such as their name, school, or city; further, they contained URLs that were copied from or based off of legitimate Web sites and began with the HTTPS protocol. In contrast, phishing emails contained no personal information; further, they contained URLs that began with HTTP and contained additional characteristics of suspicious URLs (Chatterjee & Namin, 2019): five or more dots in the domain, over 75 characters in length, presence of an IP address, the @ symbol, unusual top-level domains (e.g., ".co" rather than ".com"), and misspellings (e.g., "wwvv." rather than "www."). In sum, phishing emails differed from legitimate emails in that they contained a malicious URL and no personal information (Butavicius et al., 2015).

Emails contained one of three consequence amounts: small ($5), medium ($50), and large ($5,000). A manipulation check study (N = 70) demonstrated that requests for $5 were rated as less consequential than those for $50, which were rated as less consequential than those for $5,000.

Emails contained one of two types of consequences: gain or loss (Parsons et al., 2015). Emails that provided an opportunity to *gain* money included those in which recipients could win or save money by clicking the link. Emails that provided a threat of money *loss* included those in which recipients would need to click the link to avoid losing money, for example due to late fees or a fraudulent charge.

Emails contained one of two urgency levels: present or absent. In urgency present emails, participants were told to respond in 48 hours or less (e.g., 24 hours, midnight tonight). Urgency absent emails did not mention a response time frame.

A demographics questionnaire was also created for this study. It asked about participants' age and gender.

Procedure

Participants completed the phishing detection task and the demographic questions online using Qualtrics (2021). During the phishing detection task, participants were told to role play (Downs et al., 2006). The role play procedure allowed for the email stimuli to reference personal information about the email recipient without requiring participants to provide such information. The role participants played was of a university student "Casey Smith." This name was chosen to be both

androgynous and generic. In each of 24 trials, participants were shown one of the email stimuli. Stimulus order was randomized. The email images were presented at the top of the screen, and participants were then asked whether the email was a phishing email (Canfield et al., 2016). Participants were not provided feedback about their responses. After the phishing detection task, participants completed demographic questions concerning their age and gender.

RESULTS AND DISCUSSION

The dataset was split randomly in half, forming an initial sample (N = 177) and a validation sample (N = 176). The two datasets did not differ with regard to age or gender composition. All analyses were run twice in SPSS (Version 25), once with the initial sample and again with the validation sample. This was done to verify that results were replicable.

Overall Phishing Suspiciousness

For legitimate emails, participants correctly judged them as legitimate 66.1% (initial sample) and 67.2% (validation sample) of the time, and incorrectly judged them as phishing 33.9% (initial) and 32.8% (validation) of the time. For phishing emails, participants incorrectly judged them as legitimate 62.7% (initial) and 62.3% (validation) of the time and correctly judged them as phishing 37.3% (initial) and 37.7% (validation) of the time. A Chi Square test indicated a significant relationship between email type and phishing judgments for both samples (initial Pearson $\chi^2(1, N = 4248) = 5.320$, $p = .021$; validation Pearson $\chi^2(1, N = 4224) = 11.222$, $p = .001$). For all samples, post-hoc tests indicated participants judged emails to be legitimate more often than phishing; however, the rate at which they did so differed depending on whether the email was a legitimate or phishing email. Participants judged emails to be legitimate more often they were legitimate compared to when they were phishing. Conversely, participants judged emails to be phishing less often when they were legitimate compared to when they were phishing.

Effects of Unrealistic Promises and Urgent Wording on Phishing Suspiciousness

Binary logistic regression models were conducted to predict the users' ratings of phishing and legitimate emails as phishing or not phishing (i.e., suspiciousness). Seven variables were entered as predictors. Gender and age were entered into the model first to serve as control variables; both have previously been shown to predict phishing suspiciousness (gender: Halevi et al., 2013; 2015; Hong et al., 2013; Jagatic et al., 2007; Mohebzada et al., 2012; age: Jagatic et al., 2007, Kumaraguru et al., 2009; Parsons et al., 2019). The three manipulated variables were included in the order: urgency, consequence type, and consequence amount. Because amount had three levels, we used dummy variable coding to result in comparisons among all three levels. None of the predictors were significantly correlated with each other. The results of these regression analyses are summarized in Table 1. The following sub-sections describe results related to our hypotheses regarding unrealistic promises and urgency, and demographic predictors.

Unrealistic Promises. We hypothesized that unrealistic promises, indicated by a *gain* consequence and a *high* [$5,000] consequence amount, would serve as a phishing cue. As such, users would rate phishing emails with these characteristics as phishing more often (i.e., be more suspicious) than other phishing emails without these cues.

Supporting this hypothesis, users were more suspicious of phishing emails that had a *gain* consequence than a *loss* consequence in both the initial (Wald $\chi^2 = 191.099$, $p < .001$) and validation (Wald $\chi^2 = 144.752$, $p < .001$) models. This finding was also true for legitimate emails that had a *gain* consequence in both the initial (Wald $\chi^2 = 97.044$, $p < .001$) and validation (Wald $\chi^2 = 84.844$, $p < .001$) models. These results indicate that users attended to the *gain* cue for both phishing and legitimate emails to make their phishing judgments.

In additional support for this hypothesis, users were most suspicious of phishing emails that had a *large* [$5,000] consequence amount compared to a *small* [$5] or *medium* [$50] consequence amount in both the initial (*small* Wald $\chi^2 = 209.619$, $p < .001$; *medium* Wald $\chi^2 = 117.935$, $p < .001$) and validation (*small* Wald $\chi^2 = 216.838$, $p < .001$; *medium* Wald $\chi^2 = 118.083$, $p < .001$) models. These results indicate that users additionally attended to the *large* consequence amount when making their judgments about phishing emails.

Interestingly, for legitimate emails, users were most suspicious of the *medium* [$50] consequence amount compared to the *small* [$5] and *large* [$5,000] consequence amounts in both the initial (*small* Wald $\chi^2 = 117.331$, $p < .001$; *large* Wald $\chi^2 = 48.756$, $p < .001$) and validation (*small* Wald $\chi^2 = 96.913$, $p < .001$; *large* Wald $\chi^2 = 56.593$, $p < .001$). Legitimate emails differed from phishing emails in that they 1) mentioned personal information of the user's persona and 2) contained a non-malicious URL. In the presence of personal information and without a malicious URL cue, users may be attending to other information to determine whether an email is a phishing attack. Spear-phishing attacks, a subset of phishing attacks, aim to masquerade as legitimate emails from a trusted source compared to phishing emails, meaning they often mention more "realistic" amounts of money instead of large sums. Therefore, it may be that users are attempting to distinguish between legitimate and *spear-phishing* emails, rather than *phishing* emails, when the email includes personal information and lacks a typical phishing cue (i.e., a malicious URL).

Urgency. We hypothesized that *urgency* would serve as a phishing cue such that users would rate phishing emails with this characteristic as phishing more often (i.e., be more suspicious) than other phishing emails without these cues.

Contradicting our hypothesis, for phishing emails, urgency was not a significant predictor of suspiciousness in either the initial (Wald $\chi^2 = 0.000$, $p = 1.000$) or validation (Wald $\chi^2 = .198$, $p = .056$) model. This result indicates that for phishing emails, consequence type and amount were more important phishing cues for users than urgency.

Interestingly, users were more suspicious of legitimate emails that were *urgent* in both the initial (Wald $\chi^2 = 67.434$, $p < .001$) and validation (Wald $\chi^2 = 73.331$, $p < .001$) model. Similar to the results for unrealistic promises, this result may indicate that urgency serves as an alternative phishing cue when an email contains personal information and lacks a typical

phishing cue (i.e., a malicious URL). Users may also be attending to urgency in these situations because it is less common to receive personal emails, otherwise free of phishing cues, that express urgency.

Demographics. We investigated the effects of two demographic variables, age and gender, on phishing suspiciousness. As noted earlier, both have previously been shown to predict phishing suspiciousness (age: Jagatic et al., 2007, Kumaraguru et al., 2009; Parsons et al., 2019; gender: Halevi et al., 2013; 2015; Hong et al., 2013; Jagatic et al., 2007; Mohebzada et al., 2012).

Age did not predict phishing suspiciousness in any of the four regression models. Although age has been shown to predict phishing suspiciousness, it does not always (e.g., Mohebzada et al., 2012, Zielinska et al., 2014). As such, the present findings are consistent with those exceptions. It is important to note, however, that our sample was skewed heavily toward the 18–25-year-old range. Our age range may be too narrow to detect differences due to age, and thus we cannot draw definitive conclusions about the role of age in phishing suspiciousness.

Gender did predict phishing suspiciousness, but only in certain models. In the initial model for legitimate emails, females judged emails as phishing more often than males (Wald $\chi^2 = 7.241$, $p = .007$), indicating that females were more suspicious of legitimate emails than males. However, this result did not replicate in the validation model for legitimate emails. In the initial model for phishing emails, gender did not predict phishing suspiciousness. However, in the validation model for phishing emails, females judged phishing emails as phishing less often than males (Wald $\chi^2 = 6.621$, $p = .010$), indicating that males were more suspicious of phishing emails than females. Overall, these results suggest gender is an inconsistent predictor of phishing suspiciousness. Consistent with that, the literature supports that gender can predict phishing suspiciousness, but does not always (e.g., Kumaraguru et al., 2009, Parsons et al., 2019).

CONCLUSIONS

This study investigated whether urgent wording and unrealistic promises affected user suspiciousness of phishing and legitimate emails. We predicted that users would be more suspicious of phishing emails that were *urgent* and contained offers of *large gains* (i.e., "too good to be true"). The results partially support these predictions. Users were more suspicious of phishing emails that contained *gains* as opposed to *losses*. Users were also most suspicious of phishing emails that contained a *large* consequence amount compared to *small* and *medium* amounts. These findings suggest that unrealistic promises were an important cue for users to detect phishing emails. *Urgency*, however, only significantly predicted suspiciousness for legitimate emails, suggesting that *urgency* may only be attended to in the absence of other well-known phishing cues. These findings indicate that, depending on email type, users differently attend to the cues of urgent wording and unrealistic promises in judging whether an email is phishing.

These findings have important cybersecurity implications. For example, penetration testing emails that focus on unrealistic promises and urgency may only lure the most susceptible users; the presence of such cues will likely cue other users. Further, anti-phishing training should inform users that they often attend to unrealistic promises and urgency when judging phishing emails, and warn to be cognizant of other cues that may be inadvertently missed.

REFERENCES

Anti-Phishing Working Group (APWG). (2021). *Phishing activity trends report. 4th quarter 2020.* 1-14.

Butavicius, M., Parsons, K., Pattinson, M., & McCormac, A. (2015). Breaching the human firewall: Social engineering in phishing and spear-phishing emails. *Australasian Conf. on Information Systems 2015 Proceedings*, *98*.

Canfield, C. I., Fischhoff, B., & Davis, A. (2016). Quantifying phishing susceptibility for detection and behavior decisions. *Human Factors*, *58*(8).

Chatterjee, M., & Namin, A.S. (2019). Detecting phishing websites through deep reinforcement learning. *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, *2*, 227-232.

Cialdini, R. B. (2001). *Influence: Science and practice*. Allyn & Bacon.

CyberEdge Group. (2020). *2020 Cyberthreat defense report*. 1-58.

Downs., J. S., Holbrook, M., & Cranor, L. F. (2006). Decision strategies and susceptibility to phishing. *Symposium on Usable Privacy and Security (SOUPS)*, 79-90.

Halevi, T., Lewis, J., & Memon, N. (2013). A Pilot Study of Cyber Security and Privacy Related Behavior and Personality Traits. *WWW '13 Companion: Proceedings of the 22nd International Conference on World Wide Web*. 737-744. https://doi.org/10.2139/ssrn.2383427

Halevi, T., Memon, N., & Nov, O. (2015). Spear-Phishing in the Wild: A Real-World Study of Personality, Phishing Self-Efficacy and Vulnerability to Spear-Phishing Attacks. SSRN Electronic Journal.

Hong, K. W., Kelley, C. M., Tembe, R., Murphy-Hill, E., & 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, 57(1), 1012–1016.

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

Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M. A., & Pham, T. (2009). School of phish: A real world evaluation of anti-phishing training. *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09*.

Lötter, A., & Futcher, L. (2015). A framework to assist email users in the identification of phishing attacks. *Information and Computer Security*. *23*(4).

Mohebzada, J. G., Zarka, A. E., Bhojani, A. H., & Darwish, A. (2012). Phishing in a university community: Two large scale phishing experiments. *2012 International Conference on Innovations in Information Technology (IIT).*

Parsons, K., Butavicius, M., Delfabbro, P., & Lillie, M. (2019). Predicting susceptibility to social influence in phishing emails. *International Journal of Human-Computer Studies*, *128*, 17–26.

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

Proofpoint. (2021). *2021 State of the phish: An in-depth look at user awareness, vulnerability and resilience*. 1-43.

Qualtrics (2021). https://www.qualtrics.com

Vishwanath, A., Herath, T., Chen, R., Wang, J., & 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*(3), 576–586.

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

Zielinska, O. A., Tembe, R., Hong, K. W., Ge, X., Murphy-Hill, E., & Mayhorn, C. B. (2014). One Phish, Two Phish, How to Avoid the Internet Phish: Analysis of Training Strategies to Detect Phishing Emails. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*(1).

Zielinska, O. A., Welk, A. K., Mayhorn, C. B., & Murphy-Hill, E. (2015). Exploring expert and novice mental models of phishing. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 59(1), 1132-1136.

Table 1

Binary Logistic Regression Models Predicting Suspiciousness for Legitimate and Phishing Emails

Model	*B*	SE *B*	Wald χ^2	*OR*	95% CI for *OR*
Legitimate Initial (R^2 = .135)					
Male – Female	.329	.122	7.241**	1.389	[1.093, 1.765]
Age	-.002	.015	.012	.998	[.969, 1.028]
Not Urgent – Urgent	**.843**	**.103**	**67.434*** **	**2.324**	**[1.901, 2.843]**
Loss – Gain	**1.018**	**.103**	**97.044*** **	**2.767**	**[2.260, 3.388]**
Amount			123.209***		
Small [$5] – Med [$50]	**1.387**	**.128**	**117.331*** **	**4.002**	**[3.114, 5.144]**
Med [$50] – Large [$5000]	**-.834**	**.119**	**48.756*** **	**.434**	**[.344, .549]**
Small [$5] – Large [$5000]	**.553**	**.131**	**17.890*** **	**1.738**	**[1.345, 2.245]**
Constant	-2.579	.357	52.127***	.076	
Legitimate Validation (R^2 = .126)					
Male – Female	-.017	.118	.021	.983	[.780, 1.240]
Age	-.016	.015	1.104	.984	[.956, 1.014]
Not Urgent – Urgent	**.904**	**.105**	**74.336*** **	**2.470**	**[2.011, 3.034]**
Loss – Gain	**.965**	**.105**	**84.228*** **	**2.624**	**[2.135, 3.224]**
Amount			109.674***		
Small [$5] – Med [$50]	**1.266**	**.128**	**97.102*** **	**3.546**	**[2.757, 4.561]**
Med [$50] – Large [$5000]	**-.924**	**.123**	**56.593*** **	**.397**	**[.312, .505]**
Small [$5] – Large [$5000]	**.342**	**.133**	**6.620***	**1.408**	**[1.085, 1.827]**
Constant	-1.994	.342	34.068***	.136	
Phishing Initial (R^2 = .196)					
Male – Female	.057	.121	.220	1.059	[.834, 1.343]
Age	-.006	.016	.160	.994	[.964, 1.025]
Not Urgent – Urgent	.000	.103	.000	1.000	[.817, 1.224]
Loss – Gain	**1.490**	**.108**	**191.099*** **	**4.438**	**[3.593, 5.482]**
Amount			230.243***		
Small [$5] – Med [$50]	**.576**	**.132**	**19.207*** **	**1.780**	**[1.375, 2.303]**
Med [$50] – Large [$5000]	**1.342**	**.124**	**117.935*** **	**3.825**	**[3.003, 4.873]**
Small [$5] – Large [$5000]	**1.918**	**.132**	**209.619*** **	**6.808**	**[5.251, 8.827]**
Constant	-2.130	.361	34.788***	.119	
Phishing Validation (R^2 = .186)					
Male – Female	-.303	.118	6.621*	.739	[.586, .930]
Age	.007	.014	.231	1.007	[.979, 1.036]
Not Urgent – Urgent	.188	.104	3.276	1.206	[.985, 1.478]
Loss – Gain	**1.293**	**.107**	**145.079*** **	**3.643**	**[2.952, 4.496]**
Amount			238.540***		
Small [$5] – Med [$50]	**.614**	**.133**	**21.273*** **	**1.847**	**[1.423, 2.398]**
Med [$50] – Large [$5000]	**1.339**	**.123**	**118.083*** **	**3.816**	**[2.997, 4.858]**
Small [$5] – Large [$5000]	**1.953**	**.133**	**214.720*** **	**7.049**	**[5.429, 9.154]**
Constant	-2.142	.335	40.873***	.117	

Note. R^2 values are Cox & Snell R^2 for the overall model. Non-standardized beta values (*B*), Standard error of beta values (SE *B*), test statistic (Wald χ^2), odds ratio (*OR*), 95% confidence ratio for odds ratio (95% CI for *OR*). Comparisons are in the format [reference group] – [comparison group] (e.g., Loss – Gain compares the gain group to the loss group). Bolded lines indicate that a variable was significantly predictive of suspiciousness in *both* the initial and the validation models for that email type.
***p < .001, **p < .01, *p < .05.