Int. J. Human–Computer Studies 174 (2023) 103017

Contents lists available at ScienceDirect

International Journal of Human - Computer Studies

journal homepage: www.elsevier.com/locate/ijhcs

Phishing interrupted: The impact of task interruptions on phishing email classification

Elisabeth J.D. Slifkin, Mark B. Neider [*]

Department of Psychology, University of Central Florida, USA

ARTICLE INFO

Keywords:
Phishing
Task interruptions
Cybersecurity
Phishing emails

ABSTRACT

Email is a pervasive form of communication in both personal and professional settings. The extent to which a human user can accurately detect phishing emails impacts the amount of risk they are exposed to. Previous research has unfortunately shown that people often fall victim to phishing attacks, both in controlled laboratory and naturalistic settings, even if they have received phishing awareness training. While the impact of numerous participant, email, and task characteristics on phishing email classification performance have been assessed, the impact of various environmental factors must be considered as well. Interruptions are a widespread occurrence in daily life and have been shown to negatively impact performance on many tasks. The present studies aim to explore the effect of task interruptions on phishing email classification. Participants performed a classification task where they categorized emails as either phishing or legitimate. Participants in both experiments were occasionally interrupted with either a secondary task to complete (Experiments 1 & 2) or a blank box (Experiment 2). Results of both experiments indicated higher phishing classification accuracy on interrupted trials and an increase in response time, roughly equal to the amount of time the email was shown prior to the interruption, regardless of the type of interruption. Participants also showed an unbiased response and were more sensitive to the task when interrupted. Our findings suggest that being interrupted during this phishing classification task may improve performance in this limited capacity, though this benefit may not be due to the interruptions directly.

1. Introduction

Days often begin with a cup of coffee and checking email, as email is an essential part of everyday life, both professionally and personally. Email is constant, invasive, and frequently engaged with while completing other activities. During the workday, people may check their email up to 77 times on average, spending up to 1.5 h reading and replying to emails (Mark et al., 2016). Approximately half of workers in the United States check their personal email every few hours during work hours, as well as their work email every few hours outside of business hours (Ceci, 2022). Checking email often requires dedicated time to complete, yet sitting down and having 20 min to read and reply to emails sounds like a luxury for many, as there are often numerous interruptions that occur during these times. When interrupted during a task, it often takes more time to complete that task and more errors are committed (Bailey and Konstan, 2006). When reading email, clicking on a malicious link or replying with personal information can compromise personal security. As interruptions are disruptive to many tasks, it is reasonable to assume that they would disrupt the task of email classification as well. Thus, when not focused solely on reading emails due to being interrupted, the chance of experiencing undesirable outcomes would likely increase. In many cases, reading email is not a task that exists in a vacuum, and therefore it should not be studied in one. Situational factors, such as interruptions, can have increasingly negative effects on many tasks, and thus must be considered when examining phishing email classification.

1.1. Phishing emails

With an estimated 4.03 billion email users worldwide, approximately 306.4 billion emails were sent and received per day in 2020 (Ceci, 2022). These numbers are expected to continue rising in the next few years, and the number of phishing emails will likely also rise. Phishing is an attempt to unlawfully obtain users' personal information *via* email (Drake et al., 2004). Phishing attacks are the most common form of email account compromise (EAC) attacks and are one of the

[*] Corresponding author.
E-mail address: mark.neider@ucf.edu (M.B. Neider).

https://doi.org/10.1016/j.ijhcs.2023.103017
Received 1 July 2022; Received in revised form 14 February 2023; Accepted 26 February 2023
Available online 1 March 2023

greatest contributors to cybercrime in the United States (Gorham, 2020). Over 450,000 internet crime complaints, totaling an estimated $3.5 billion in losses from both individuals and corporations, were reported to the Federal Bureau of Investigation (FBI) in 2020, with $1.7 billion of these losses directly attributed to EAC attacks (Gorham, 2020).

Phishing emails often mimic reputable companies, seek to create a plausible premise, contain numerous spelling or grammatical errors, require an immediate response, collect personal information directly in the email, and over-emphasize security (Drake et al., 2004). Increasing cybersecurity awareness and bolstering email protocols are important steps in reducing the dangers posed by phishing emails, however they will not eliminate the threat entirely. Using various machine learning techniques, some email filters can detect and intercept up to 99% of phishing emails, with limited numbers of legitimate emails being caught incorrectly (Bergholz et al., 2010; Gangavarapu et al., 2020). This is an impressive statistic, but it also indicates that a minimum of 1% of phishing emails still make their way into users' inboxes. For a user who receives 1000 emails per week, 10 of those will be phishing. Unless email filters can reliably detect 100% of phishing emails, the end user retains some level of responsibility for determining the legitimacy of the emails in their inbox, making the human factor a critical piece toward achieving maximal security against email-centered security attacks. The extent to which a human user can accurately detect phishing emails impacts the amount of risk they are exposed to. The less reliably someone can identify phishing emails, the more at risk they are for falling victim to a phishing attack. Unfortunately, research has shown that most human end users are not able to reliably determine the legitimacy of emails with high levels of accuracy.

Analyzing users' email behavior is generally conducted through online surveys, naturalistic studies, or laboratory experiments. Surveys ask users to self-report previous phishing experience and aim to identify potential factors for phishing vulnerability (Grimes et al., 2007; Sheng et al., 2010). Both naturalistic and laboratory studies of phishing utilize simulated phishing emails, where these emails are either manipulated by the experimenters or are taken from a set of previously identified phishing emails. Naturalistic studies frequently send fake phishing emails to users' personal, school, or business email accounts and subsequently collect data on any interactions with those emails (Oliveira et al., 2017; Vishwanath et al., 2011, 2018). Laboratory experiments often ask participants to classify emails from a test set and may ask them to select the action they would then take with that email (e.g., reply, delete, ignore; Canfield et al. 2016, Nyeste and Mayhorn 2010, Sarno et al. 2017, 2020). In both naturalistic and laboratory settings, large percentages of participants often fall victim to the simulated phishing attacks or fail to correctly classify email as legitimate or illegitimate. Previous research has shown that up to 47% of participants fail to identify or misclassify phishing emails, with up to 28% doing so even after training (Canfield, 2016; Sarno et al., 2020; Sheng, 2010).

While previous phishing research has provided evidence that there are combinations of participant (Grimes et al., 2007; Hong et al., 2013; Kumaraguru et al., 2007; Li et al., 2020; Mayhorn and Nyeste, 2012; Olivera et al., 2017; Sarno et al., 2017, 2020; Sheng et al., 2010), email (Drake et al., 2004; Patel et al., 2019; Williams et al., 2019), and task characteristics (Sarno et al., 2017, 2020, 2022) that may impact phishing susceptibility, factors associated with the external environment have been less of a focal point. Even in laboratory settings where controlling the external environment to minimize task-irrelevant distractions is relatively easy, people are unable to reliably classify emails as phishing. This implies that other factors may affect phishing email identification. In the real world, people check email in a variety of environments that may influence their current level of susceptibility to phishing emails.

The process of detecting phishing emails is a difficult, but important undertaking; people are tasked with determining the legitimacy of the content in front of them, with potentially severe consequences if not performed correctly. Since the goal of a phishing attack is to obtain personal information, the consequences are similar to identity theft, including financial, physical, and emotional distress. While most individuals suffer minor financial loss (less than $100), some experience losses greater than $20,000 (ITRC, 2021; Li et al., 2019). Physically, people may have trouble sleeping, experience headaches, upset stomach, fatigue, or high blood pressure (Golladay and Holtfreter, 2017). Many people feel worried, anxious, depressed, vulnerable, and violated, among other things, after being victims of identity theft (Golladay and Holtfreter, 2017). While most identity theft cases can be resolved in about a month, some take more than a year to settle (ITRC, 2021). Falling victim to a phishing attack and experiencing identity theft can be life altering events, so it is important to understand how to identify phishing emails and protect oneself from potential attacks.

Considering participants are unable to consistently detect phishing emails, missing up to 47% of phishing emails in test sets even in a controlled laboratory environment, it is relevant to consider the role that the external environment may play in the completion of this task. One aspect of the external environment that has been shown to greatly impact performance in numerous contexts and tasks is interruptions. As interruptions are widespread threats to productivity, it is important to determine the impact that interruptions have on one's ability to correctly, and reliably, detect phishing emails.

1.2. Task interruptions

Interruptions are an unavoidable consequence of working in the presence of other individuals and copious amounts of technology. The increase in technology has multiplied the ways in which, and the ease by which, people can be interrupted. In modern western society, most people can be reached with relative ease *via* cell phone (calling or texting), smartwatch, office phone, email, instant message, or video conference. Eighty-five percent of Americans owned a smartphone as of 2021, with about a third of individuals reporting being online and reachable almost constantly (Pew Research Center, 2021). The average office worker is now interrupted three to 10 times per day, with some individuals experiencing upwards of 20 interruptions (Leroy and Glomb, 2018). An interruption occurs when we divert our attention toward something distracting and usually unanticipated, often at the expense of our productivity on the original task (Couffe and Michael, 2017). Though the interruption literature spans numerous disciplines (e.g., human-computer interaction, information technology, healthcare), it lacks a unified definition (see Puranik et al. 2020 for review). An interruption can be most broadly defined as an unexpected intrusion that requires the suspension of a task's execution, with or without the suspension of attentional focus (Puranik et al., 2020). Interruptions often also include the intention of returning to the original task (Altmann and Trafton, 2002; Boehm-Davis and Remington, 2009; Couffe and Michael, 2017).

Interruptions come in many forms and can be disruptive to ongoing tasks. The effect of interruptions on task performance can be quantified broadly in terms of time costs and error rates. Interruptions seem to reliably increase resumption time and/or total time on task (Bailey and Konstan, 2006; Hodgetts and Jones, 2006), however accuracy is less reliably affected. Sometimes it is impaired, as evidenced by an increase in error rates (Altmann et al., 2014; Foroughi et al., 2015), and other times it is not (Williams and Drew, 2017). People who are interrupted during a task may fail to return to the original task in a timely manner, or may neglect the original task altogether if not prompted to continue (O'Conaill and Frolich, 1995). Individuals may also require up to 27% more time to complete the main task when interrupted (Bailey and Konstan, 2006). Not all interruptions affect performance equally, however; more frequent interruptions and those unrelated to the main task result in decreased decision accuracy and increased decision time on the main task (Speier et al., 1999). Additionally, the extent to which an interruption visually occludes the main task impacts resumption ability. The visibility of the main task during an interruption is associated with faster and more accurate recovery (Iqbal and Horvitz, 2007; Ratwani and Trafton, 2008; Ratwani et al., 2007), while interruptions that fully occlude the main task have shown to be more disruptive than those that only partially occlude the main task (Ratwani et al., 2007). On a cell phone, a banner text message notification is representative of a partial occlusion, while a full-screen phone call notification is typical of a fully occlusive interruption. Over time, individuals can adapt to particularly disruptive interruptions, however. While providing a warning prior to an interruption (e.g., a phone ringing) often allows for faster resumption of the main task than when no warning is present (e.g., a popup that suddenly occludes the entire display), individuals will begin to resume the main task faster after numerous interruptions without warning, indicating adaptation to the interruption (Trafton et al., 2003).

1.2.1. *Memory for goals*

Altmann and Trafton's (2002) memory for goals theory can be used to explain why interruptions are disruptive by contextualizing the effects of interruptions in terms of memory activation and associative priming. Essentially, a primary task goal is activated, but when an interruption is present, the interrupting task goal overwrites the primary task goal. This causes the primary task goal to be suspended and begin to decay. The longer an interruption is, the more the primary task goal decays, thereby increasing the time it takes to recall and return to the primary task. After an interruption disappears, the original task goal must be reactivated prior to resumption. This reactivation process causes the primary task goal to overwrite the interrupting task goal (Altmann and Trafton, 2002). Without this reactivation however, the original task will not be resumed.

1.3. *The present studies*

In laboratory settings where interruptions and distractions can be reduced to near-zero, people still struggle to reliably identify phishing emails. Outside laboratory settings, individuals have demonstrated to be no better at correctly classifying emails as phishing. In the real world, interruptions and distractions occur frequently, potentially contributing to poor task performance. While previous literature on phishing classification considers numerous factors that influence performance, the impact of environmental characteristics, such as interruptions, is overlooked. To successfully advance the research on phishing classification, these environmental factors must be considered.

While interruptions have often been studied in the context of sequential computer tasks with which participants have limited or no prior experience, Williams and Drew (2017) assessed the impact of interruptions on diagnostic radiology, a less systematic task where participants were quite familiar with the requirements. In this study, they found a cost of interruptions on time, but no accuracy decrement. This was attributed to impaired memory for previously searched regions of images, as evidenced by refixations (Williams and Drew, 2017). Additionally, Foroughi and colleagues (2015) evaluated the effect of interruptions on reading comprehension ability and found that interruptions significantly disrupted comprehension, particularly when information must be connected and synthesized across a passage. This effect can be mitigated however, by allowing for uninterrupted processing of information prior to an interruption.

In the current studies, we therefore aim to examine the impact of task interruptions on phishing email classification. Overall, we expect that interruptions will lead to a decrease in classification accuracy and an increase in classification time. The decrease in accuracy is expected since email classification is a reading task that requires the synthesis of information (Foroughi et al., 2015). Classifying emails is a task that participants have had previous experience with, thus the increase in classification time is expected, as even practiced tasks show a time cost of interruptions (Williams and Drew, 2017). Experiment 1 explored the overarching effect of interruptions on phishing email classification performance and Experiment 2 investigated the effect of interruption complexity.

2. Experiment 1

The goal of Experiment 1 was to determine the effects of task-irrelevant interruptions on email classification performance. Interruptions unrelated to the main task have been shown to decrease decision accuracy and increase decision time on the main task (Speier et al., 1999). Mathematical operations have been reliably used as interrupting tasks in previous research (Gillie and Broadbent, 1989; Hodgetts and Jones, 2006), and in this case were dissimilar to the main email task. As interruptions seem to reliably engender a time cost, even on practiced tasks (Williams and Drew, 2017), an increase in classification time is expected on trials where participants are interrupted. As the email is immediately reintroduced after completion of the interruption, the primary task goal should be reactivated and the main classification task should be resumed (Altmann and Trafton, 2002). Without location-specific priming however, participants may not remember where they stopped reading prior to the interruption and therefore may fail to pick up where they left off. Thus, it is expected that participants will take longer to classify an email when they are interrupted.

Previous research on email classification showed that participants took approximately 10 s to determine the legitimacy of an email (Sarno et al., 2020, 2022). Oulasvirta and Saariluoma (2006) found that people can resume a task with little disruption after an interruption when there is enough time to encode the information prior to being interrupted. By interrupting participants much sooner than 10 s, it is likely that they will not have had time to encode all the information in any given email. Additionally, the information presented in each email is required to be synthesized and connected prior to forming a judgment about that email, so the presence of an early-onset interruption should disrupt this process, therefore decreasing classification accuracy on interrupted trials (Foroughi et al., 2015).

2.1. *Method*

2.1.1. *Participants*

A total of 169 participants from the University of Central Florida completed this study online in exchange for partial course credit. The popup interruptions were the main task manipulation; participants who did not view more than 50% of the interruptions were excluded from analyses. If participants viewed more than 50% of the total number of interruptions, individual trials on which participants failed to view an interruption were excluded from analyses. Eighteen participants were excluded from the analyses; 15 participants were excluded for failing to view more than half of the interruptions and three were excluded for having average response times more than three standard deviations from the group's mean. The final sample consisted of 151 participants (58.3% female, $M_{age} = 19.4$ years). All participants self-reported normal or corrected-to-normal vision and normal color vision. This research was approved by the Institutional Review Board at the University of Central Florida.

2.1.2. *Study design*

The experiment consisted of a 2 (email type) x 2 (interruption presence) within-subjects design. Email type consisted of phishing and legitimate emails, and interruption presence was identified as either present or absent. On each trial, participants classified an email as legitimate or not legitimate; 50% of trials contained phishing emails and 50% contained legitimate emails. On 20% of the trials, participants were interrupted 3 s after trial onset. Half of the interruptions occurred on phishing email trials and half occurred on legitimate email trials. The interruption occluded the email message almost entirely to ensure that it would not go unnoticed and to ensure that the main task could not be continued while the interruption was present on-screen (Adamczyk and

Bailey, 2004; Iqbal and Horvitz, 2007; Ratwani and Trafton, 2008). The interruption required participants to respond to a math equation (Fig. 1b). All equations consisted of two-digit addition, using numbers between 10 and 19, with no carrying required (e.g., 16 + 12). Interrupted trials were evenly distributed between phishing and legitimate emails. Participants completed 100 trials distributed across five experimental blocks. All conditions were randomized within blocks. No practice trials were given.

2.1.3. Stimuli and procedure

This experiment was conducted online *via* PsychoPy's Pavlovia platform (Pierce et al., 2019). Participants completed the experiment on their personal computer (i.e., a desktop or laptop), at a time, and in a location of their choosing. The experimental task consisted of 100 real emails, previously used by Sarno and colleagues (2020, 2022). Emails in this set were obtained from the researchers' inboxes or web searches, with phishing and legitimate emails matched in content (e.g., a phishing email from Amazon and a legitimate email from Amazon). Use of this email set allowed for better control over the similarity of phishing and legitimate emails. Thus, any differences between the two could be better attributed to possible phishing cues and not extraneous information. Phishing emails contained at least one of the cues identified by Drake and colleagues (2004): mimicking reputable companies, plausibility of premise, numerous grammatical or spelling errors, requiring an immediate response, collecting personal information directly in the email, and overemphasizing security. Other cues may have been present in phishing emails as well, such as incorrect company logos, however participants were not informed of any phishing cues to look for either before or during the experiment. Emails were overlayed on a Gmail interface (Fig. 1a). The email set was representative of a typical inbox, including emails with banking information, social media, personal correspondence, and advertisements (see Sarno et al. 2020 for more details).On interrupted trials, participants were required to answer the math question by typing their response. Interruptions were self-paced, and submission of an answer closed the interruption window and returned the participant to the email display. Participants were unable to respond to the email while the interruption was on-screen. Participants viewed each email for as long as they liked, but if they responded to the email task within three seconds on an interruption trial, the interruption did not appear, and they simply continued to the next trial after responding to the email.

Participants provided informed consent before viewing any study-related material. A brief demographics questionnaire was completed prior to the experimental task. For each email, participants were instructed to provide a keyboard response indicating whether the email was phishing or legitimate. Participants were instructed to respond to the emails quickly, while maintaining accuracy. No feedback was provided after any response.

2.2. Results

To compare task performance across email types and interruption presence, two-way repeated measures analyses of variance (ANOVA; email type: phishing vs. legitimate; interruption presence: present vs. absent) were conducted to assess email classification accuracy and response time. Signal detection measures were also derived across interruption presence.

2.2.1. Manipulation check

The interrupting math task was the main task manipulation in this experiment, so it was important to know whether participants actively engaged with the popup interruptions. Overall, participants responded correctly to 98.3% of the interruptions ($SD = 0.03$), and the interruptions lasted 4.47 s on average ($SD = 1.81$). Results of repeated measures ANOVAs indicated that interruption response accuracy, which was high, did not differ between phishing ($M = 0.98$, $SD = 0.05$) and legitimate emails ($M = 0.99$, $SD = 0.04$; $F(1, 150) = 3.03$, $p = 0.084$, $\eta_p^2 = .02$). Interruption response time also did not differ between phishing ($M = 4.61$, $SD = 1.67$) and legitimate emails ($M = 4.34$, $SD = 2.48$; $F(1, 150) = 2.31$, $p = 0.131$, $\eta_p^2 = .02$). This confirms that participants actively engaged with the interrupting math task during the email classification task, and there was no evidence that email characteristics (i.e. type of email) impacted the completion of the interrupting task.

2.2.2. Accuracy

To determine the effect of email type and interruption presence on email classification accuracy, a two-way repeated measures ANOVA was conducted (see Fig. 2). The analysis indicated a main effect of email type ($F(1, 150) = 13.09$, $p < .001$, $\eta_p^2 = 0.08$), such that accuracy was approximately 8% lower for phishing emails than for legitimate emails. There was also a main effect of interruption presence on email classification accuracy ($F(1, 150) = 4.11$, $p = 0.044$, $\eta_p^2 = 0.03$), such that interrupted trials had 2% higher accuracy than non-interrupted trials. A significant interaction between email type and interruption presence was also noted ($F(1, 150) = 7.70$, $p = 0.006$, $\eta_p^2 = 0.05$). Bonferroni adjusted post hoc comparisons indicated that phishing accuracy

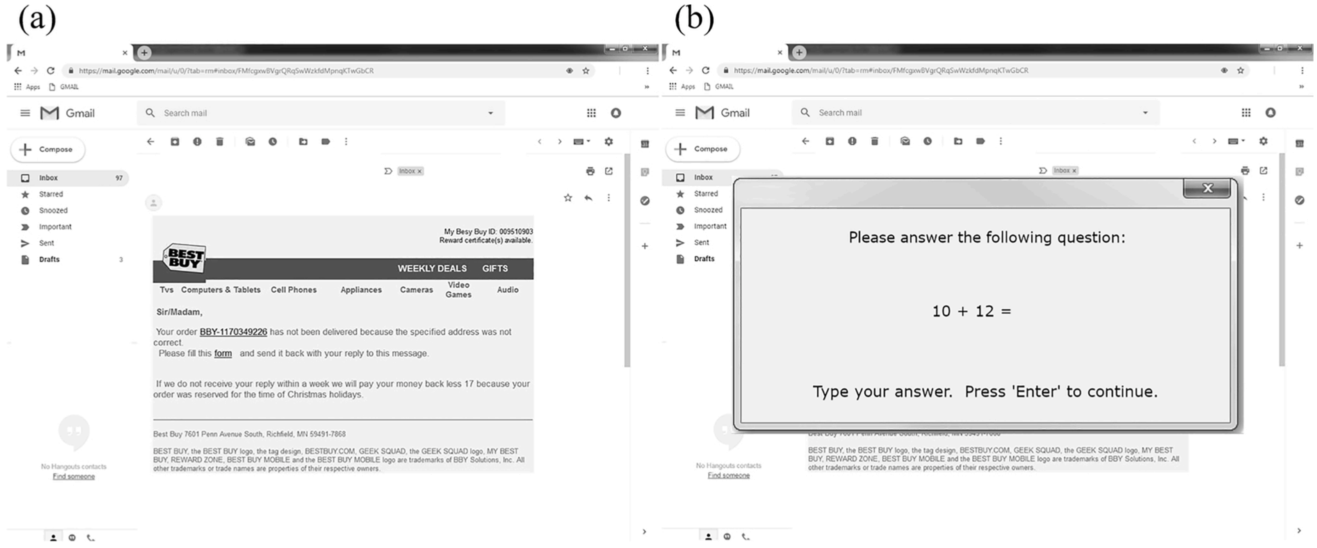

My Besy Buy ID: 009510903
Reward certificate(s) available.

BEST BUY

WEEKLY DEALS GIFTS

Tvs Computers & Tablets Cell Phones Appliances Cameras Video Games Audio

Sir/Madam,

Your order BBY-1170349226 has not been delivered because the specified address was not correct.
Please fill this form and send it back with your reply to this message.

If we do not receive your reply within a week we will pay your money back less 17 because your order was reserved for the time of Christmas holidays.

Best Buy 7601 Penn Avenue South, Richfield, MN 59491-7868

BEST BUY, the BEST BUY logo, the tag design, BESTBUY.COM, GEEK SQUAD, the GEEK SQUAD logo, MY BEST BUY, REWARD ZONE, BEST BUY MOBILE and the BEST BUY MOBILE logo are trademarks of BBY Solutions, Inc. All other trademarks or trade names are properties of their respective owners.

Please answer the following question:

10 + 12 =

Type your answer. Press 'Enter' to continue.

Fig. 1. Sample email view (a) without and (b) with an overlayed interruption.

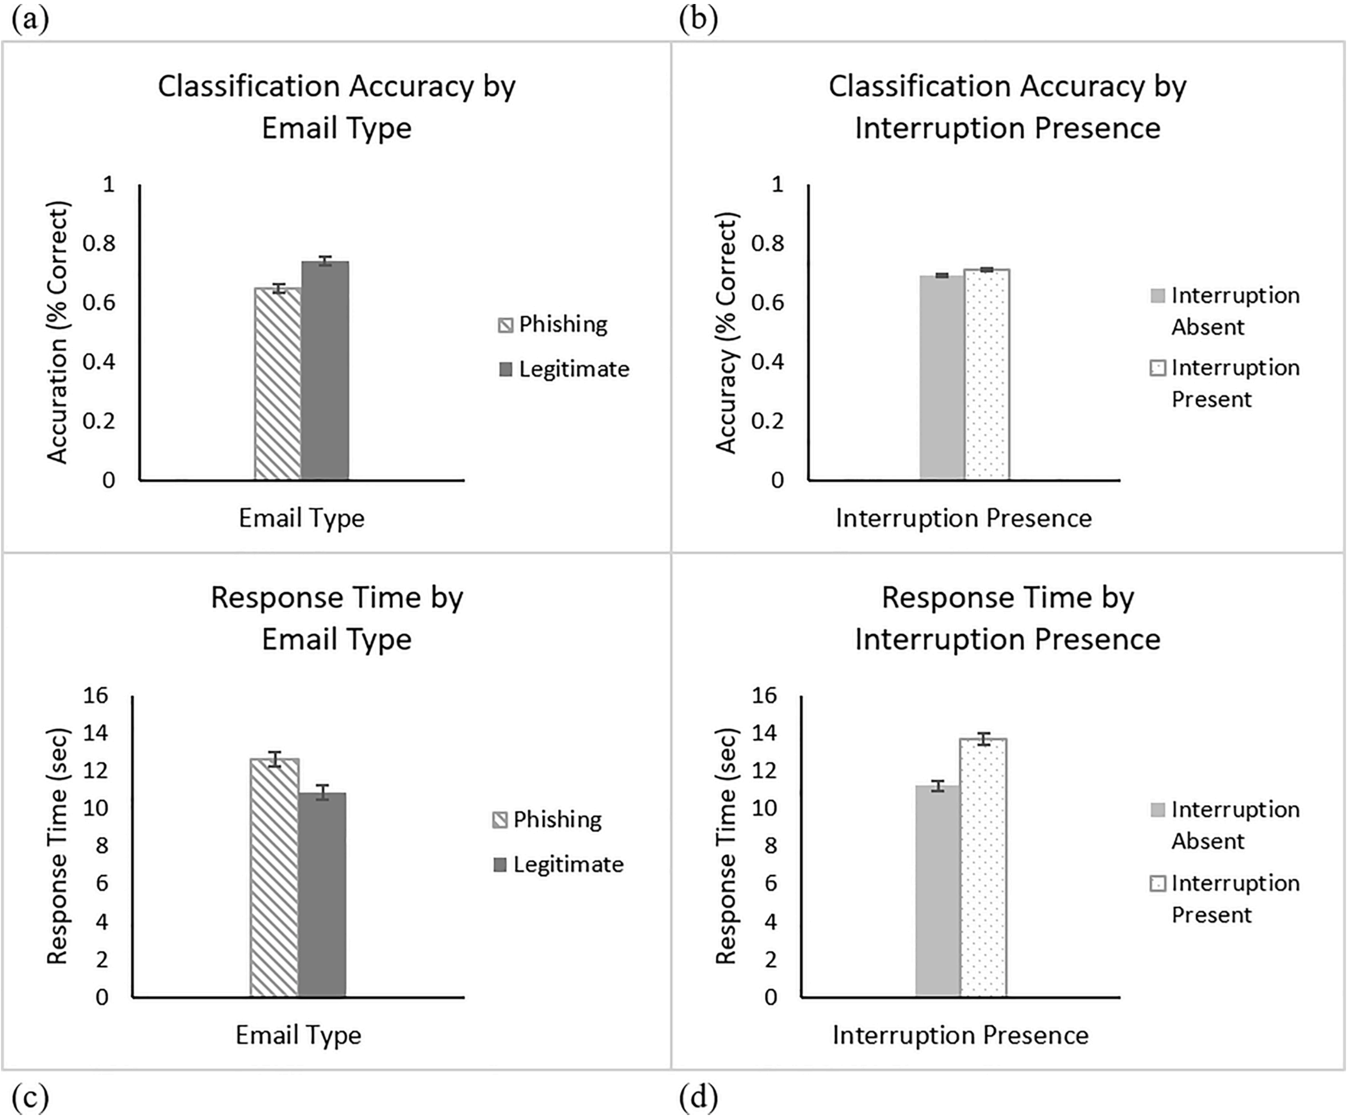

Fig. 2. Mean accuracy by (a) email type and (b) interruption presence. Mean response time by (c) email type and (d) interruption presence. Error bars indicate one standard error of the mean using the Cousineau-Morey correction method (Morey, 2008).

increased on interrupted trials (p = 0.001), but legitimate accuracy did not change based on interruption presence (p = 0.415).

2.2.3. *Response time*

To determine the effect of email type and interruption presence on email classification response time (correct trials only), a two-way repeated measures ANOVA was conducted (see Fig. 2). Response time for interruption present trials includes the 3 s of email viewing prior to the interruption appearing and the amount of time the email was viewed after the interruption ended; the time the interruption was on the screen is not included in response time.

The analysis indicated a significant effect of email type on response time ($F(1, 149) = 4.41$, $p = 0.037$, $\eta_p^2 = .03$); participants took longer (1.33 s) to respond to phishing emails than legitimate emails. Response time was also impacted by interruption presence ($F(1, 149) = 16.00$, $p < .001$, $\eta_p^2 = 0.10$), such that participants took 2.06 s longer to respond on trials where they were interrupted. There was no significant interaction between email type and interruption presence for response time ($F(1, 149) = 0.01$, $p = 0.921$, $\eta_p^2 = 0.00007$), indicating that interruption presence did not differentially impact response time on phishing and legitimate emails.

2.2.4. *Signal detection*

Signal detection measures of response bias (β) and sensitivity (d') were analyzed to elaborate on the impact of interruptions on email classification and quantify any differences in participants' inherent response tendencies due to the presence of interruptions (Green and Swets, 1966/1988). For this experiment, a hit was defined as a participant correctly identifying a phishing email; a false alarm was incorrectly identifying a legitimate email as phishing. As it relates to this experiment, response bias scores above 1 were considered more conservative in classifying an email as phishing, while scores below 1 were considered more liberal in classifying an email as phishing (Green and Swets, 1966/1988). One-way repeated measures ANOVAs were conducted to determine if there were any interruption-related differences in these measures.

There was a significant difference in response bias between interrupted and non-interrupted trials ($F(1, 150) = 7.76$, $p = 0.006$, $\eta_p^2 = 0.05$; see Fig. 3). On non-interrupted trials, a one-sample t-test indicated that participants were more conservative in their judgements, and therefore less likely to indicate that an email was phishing ($t(150) = 3.41$, $p < 0.001$, $d = 0.28$). On interrupted trials however, a one-sample t-test indicated that participants' scores were not significantly different from 1, indicating unbiased responding ($t(150) = 1.69$, $p = 0.092$, $d =$

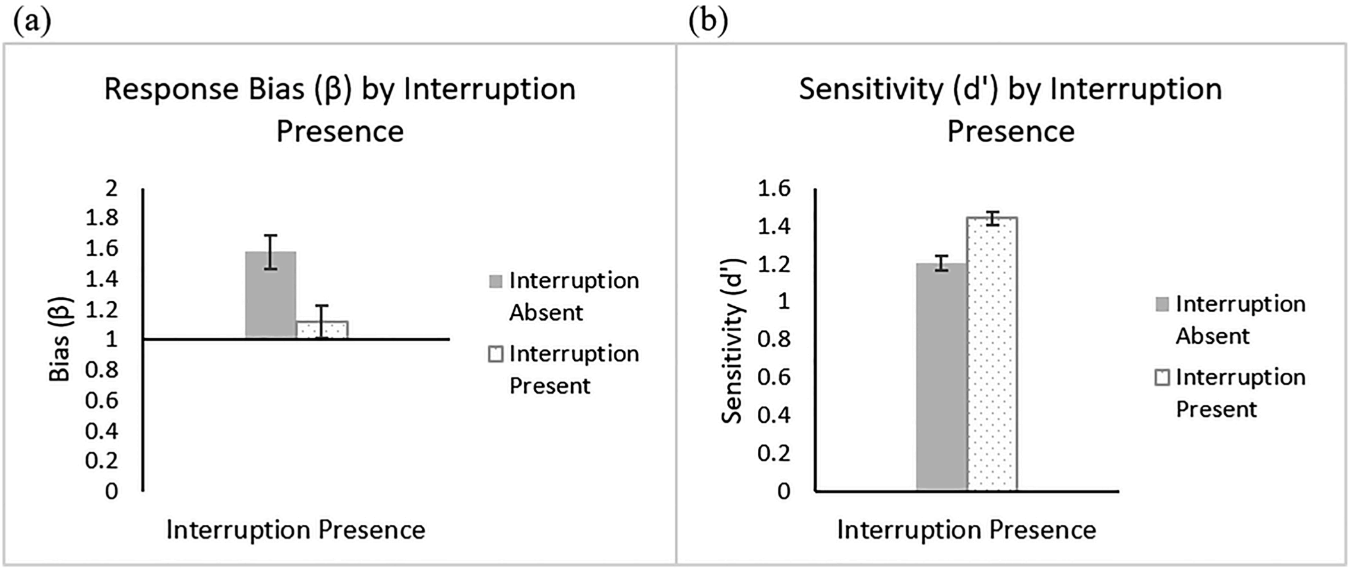

Fig. 3. Mean (a) response bias and (b) sensitivity measures by interruption presence. Error bars indicate one standard error of the mean using the Cousineau–Morey correction method (Morey, 2005).

0.14). Sensitivity (d') was also significantly different between trial types ($F(1, 150) = 19.35$, $p < 0.001$, $\eta_p^2 = 0.11$); participants were more sensitive toward distinguishing differences between phishing and legitimate emails on interrupted trials.

2.3. Experiment 1 summary

The results of Experiment 1 indicate a benefit of interruptions on phishing classification accuracy, but a cost of interruptions on response times. Participants had higher classification accuracy on legitimate emails and on interrupted trials. Interruptions served to increase accuracy on phishing emails however, while providing no benefit for legitimate emails. Interrupted trials took approximately 2 s longer than non-interrupted trials. While participants were more conservative in their judgments and therefore less likely to indicate any given email was phishing on non-interrupted trials, they were statistically unbiased on interrupted trials. This was coupled with an increase in sensitivity; when interrupted, participants were better able to distinguish between phishing and legitimate emails. The increase in phishing classification accuracy on interrupted trials, with no change in accuracy for legitimate emails, may be explained by the increase in sensitivity and the shift to unbiased responding. Participants may have become more suspicious of the emails on interrupted trials, thus creating a shift in expectation, followed by a more liberal response to classify an interrupted email as phishing.

3. Experiment 2

Experiment 1 revealed that interacting with an interruption during an email classification task improved phishing email classification accuracy. This increase in accuracy was coupled with an increase in response time when interrupted. Somewhat surprisingly, these results indicate that responding to an interruption during an email classification task serves to improve classification performance on the task, but at the expense of time spent on the task. How might this pattern of data be contextualized within the broader perspective of email classification? The increase in response time on interrupted trials in Experiment 1 was similar to the amount of time that participants viewed the email prior to being interrupted (3 s). It is possible that any judgments formed about the emails prior to the interruption were overwritten during the interruption period, causing participants to essentially reread the email from the beginning after the interruption ended; participants may have read each email, or at least portions of them, twice on interrupted trials before making a decision about legitimacy. This behavior serves to reactivate the primary task goal after the interruption (Altmann and Trafton, 2002). If the type or difficulty of the interruption differentially impacts the amount of decay of the primary task goal, and thus the need to restart the primary task, utilizing interrupting tasks of varying complexity should produce different results (Altmann and Trafton, 2002). Therefore, Experiment 2 focuses on evaluating the effect of interruption complexity on phishing email classification. Specifically, participants viewed either active (i.e., popups they had to interact with) or passive (i.e., popups they did not have to interact with) interruptions.

The presence of the interrupting math task in Experiment 1 required participants to abandon the email task and complete a different task, likely resulting in the decay of the primary (email) task goal during the interruption period (Altmann and Trafton, 2002). Presenting an interruption period without an additional task would not create a second task goal, allowing the primary task goal to remain active during the interruption period and decay more slowly, based on the length of the interruption period (Altmann and Trafton, 2002). This assumption that the primary task goal remains active during a blank interruption period means that participants should better remember where they left off on the email task prior to the interruption and therefore should resume the task more quickly and closer in proximity to where they left off. Furthermore, participants may use the passive interruption time to synthesize the information obtained from the email, allowing for a faster judgment of legitimacy. If the above are true, then participants viewing the passive interruption may not need to start the main task over after being interrupted. Thus, they should have shorter response times on interrupted trials than those viewing the active interruption. Additionally, by maintaining more of the primary task goal in memory during the interruption, participants viewing the passive interruption should have lower accuracy than those viewing the active interruption. This is because they will be better able to resume where they left off on each email, but may use their somewhat degraded memory of the primary task goal, instead of reactivating that goal by rereading the email.

3.1. Method

3.1.1. Participants

A total of 381 participants from the University of Central Florida completed this study online in exchange for partial course credit. The popup interruptions were the main task manipulation, so participants who did not view more than half of the interruptions were excluded from analyses. If participants viewed more than half of the total number of interruptions, individual trials on which participants failed to view an interruption were excluded from analyses. Sixty-five participants were excluded from analyses for various reasons; 62 participants were excluded for failing to view more than half of the interruptions and four were excluded for having average response times more than three standard deviations from the group's mean. The final sample consisted of 316 participants (58.2% female, $M_{age} = 20.1$ years). Participants were randomly assigned into active ($N = 159$) and passive ($N = 157$) interruption conditions. All participants self-reported normal or corrected-to-normal vision and normal color vision.

3.1.2. *Study design*

The experiment consisted of a 2 (interruption type) x 2 (email type) x 2 (interruption presence) mixed design with email type and interruption presence as within-subject factors and interruption type as a between-subject factor. Email type consisted of phishing and legitimate emails, and interruption presence was identified as either present or absent. All other aspects of this experiment were identical to Experiment 1, with the following exception below.

3.1.3. *Stimuli and procedure*

To assess the impact of different types of interruptions on performance, participants were divided into two interruption types: active interruptions and passive interruptions. Those in the active interruption condition viewed interruptions identical to those in Experiment 1 (Fig. 4a). Those in the passive interruption condition viewed a blank interruption box (Fig. 4b) with no arithmetic task. The active interruption continued to be self-paced, while the passive interruption remained visible for 4.5 s and then disappeared without participant input. The length of time that the passive interruption remained visible was based on the average time taken to respond to the interruption in Experiment 1.

3.2. *Results*

3.2.1. *Manipulation check*

The interrupting math task was again the main task manipulation in this experiment, so it was important to know whether participants actively engaged with the popup interruptions. Overall, participants in the active interruption condition responded correctly to 97.5% of interruptions ($SD = 0.07$) and the interruptions lasted 4.22 s on average ($SD = 1.69$). Results of repeated measures ANOVAs indicated that interruption response accuracy did not differ between phishing ($M = 0.97$, $SD = 0.08$) and legitimate emails ($M = 0.98$, $SD = 0.07$; $F(1, 158) = 0.08$, $p = 0.772$, $\eta_p^2 = 0.0005$). Interruption response time did significantly differ between phishing ($M = 4.39$, $SD = 1.72$) and legitimate emails, however ($M = 4.05$, $SD = 2.00$; $F(1, 158) = 7.00$, $p = 0.009$, $\eta_p^2 = 0.04$). Participants in the active interruption condition took longer to respond to the interruption on phishing trials. The average 4.22 s that participants in the active condition viewed the interruptions was significantly shorter than the controlled 4.5 s that participants in the passive condition viewed the interruptions ($t(158) = -2.09$, $p = 0.038$, $d = -0.17$). Despite this small difference in interruption time, no differences between groups in email classification accuracy or response time were observed, as outlined below.

3.2.2. *Accuracy*

To determine the effect of interruption type, email type, and interruption presence on email classification accuracy, a mixed ANOVA was conducted (see Fig. 5a and b). All post hoc comparisons utilized a Bonferroni correction. The analysis indicated no significant difference between the active and passive conditions on email classification accuracy ($F(1, 314) = 0.06$, $p = 0.807$, $\eta_p^2 = 0.0002$), such that participants viewing both interruption types achieved similar levels of overall accuracy (active: $M = 0.706$, $SD = 0.09$; passive: $M = 0.709$, $SD = 0.08$). There was no overall main effect of email type ($F(1, 314) = 1.92$, $p = 0.167$, $\eta_p^2 = 0.01$) or interruption presence ($F(1, 314) = 3.70$, $p = 0.055$, $\eta_p^2 = 0.01$); classification accuracy was similar for phishing and legitimate emails, as well as for interrupted and non-interrupted trials (~71%). There was a significant interaction between email type and interruption presence, however ($F(1, 314) = 14.22$, $p < 0.001$, $\eta_p^2 = 0.04$). Classification accuracy on phishing emails was higher on interrupted trials than on non-interrupted trials ($p < 0.001$), but accuracy on legitimate emails did not change ($p = 0.099$). There was no significant interaction between interruption presence and interruption type ($F(1, 314) = 0.01$, $p = 0.916$, $\eta_p^2 = 0.00004$). Classification accuracy was comparable on interrupted trials for those in the active ($M = 0.715$, $SD = 0.12$) and passive conditions ($M = 0.718$, $SD = 0.12$). A significant interaction was found between email type and interruption type, however ($F(1, 314) = 13.13$, $p < 0.001$, $\eta_p^2 = 0.04$). Participants in the active interruption condition had higher accuracy on legitimate emails compared to phishing emails ($p < 0.001$), while participants in the

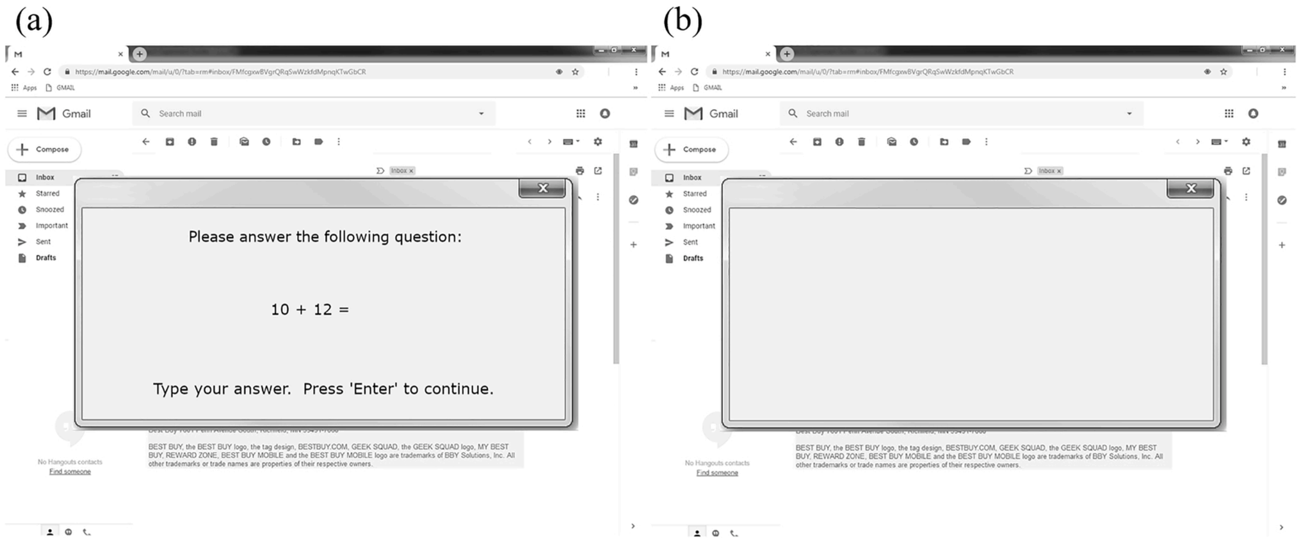

Please answer the following question:

10 + 12 =

Type your answer. Press 'Enter' to continue.

BEST BUY, the BEST BUY logo, the tag design, BESTBUY.COM, GEEK SQUAD, the GEEK SQUAD logo, MY BEST BUY, REWARD ZONE, BEST BUY MOBILE and the BEST BUY MOBILE logo are trademarks of BBY Solutions, Inc. All other trademarks or trade names are properties of their respective owners.

Fig. 4. Sample email view of (a) active and (b) passive interruptions.

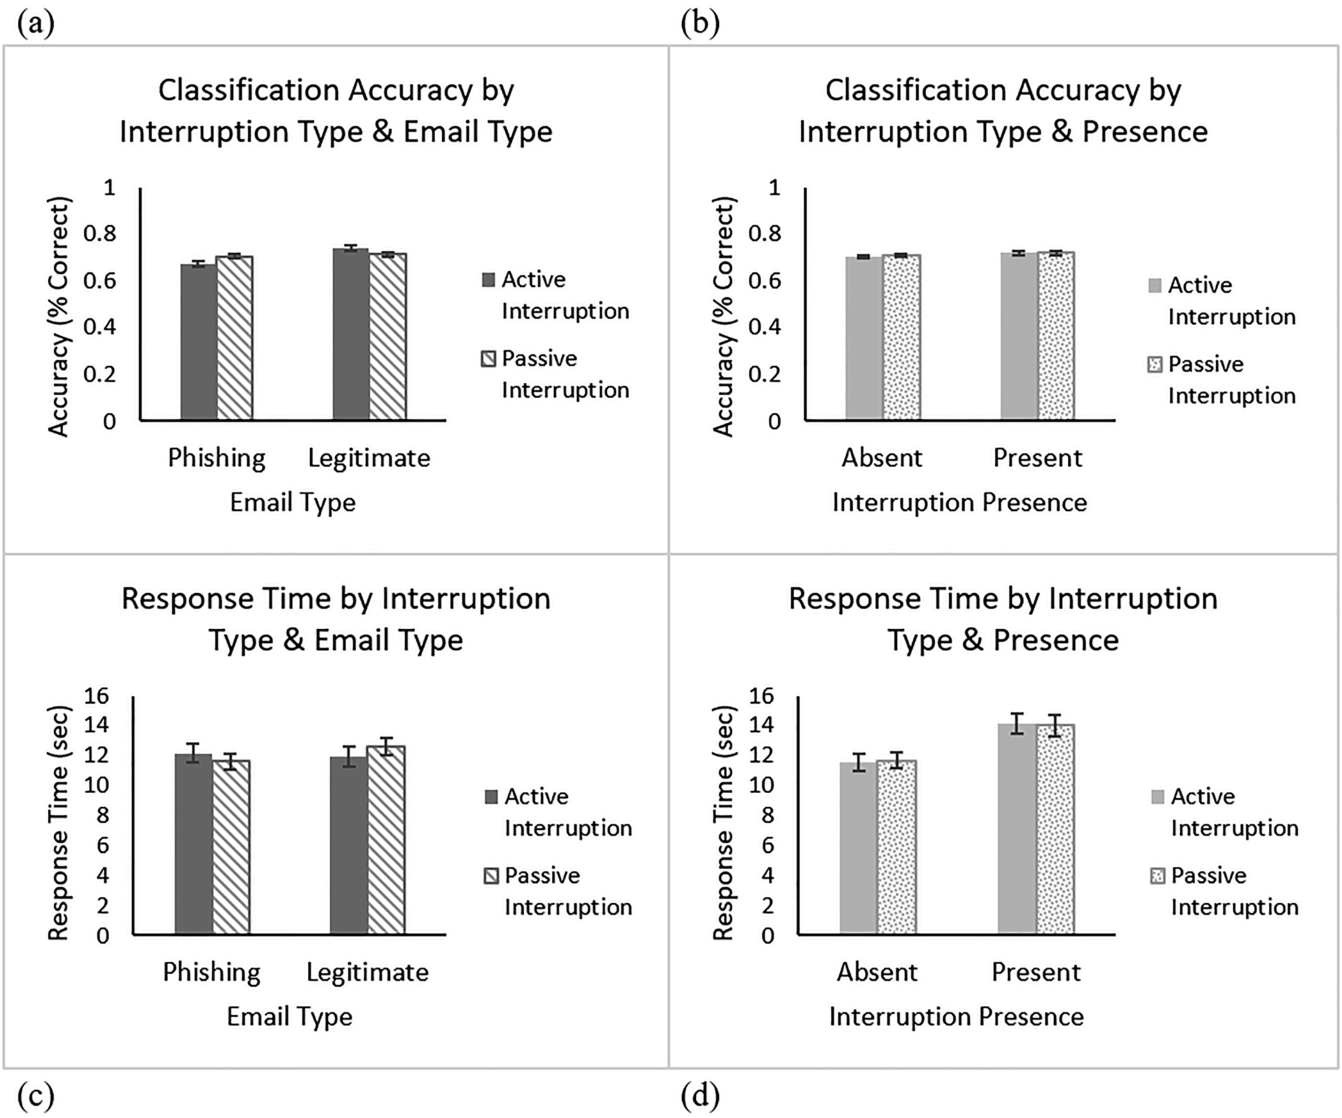

Fig. 5. Mean accuracy per interruption type by (a) email type and (b) interruption presence. Mean response time by (c) email type and (d) interruption presence. Error bars indicate one standard error of the mean.

passive interruption condition showed no difference in accuracy on phishing and legitimate emails (p = 0.116). This can be further explained by the significant three-way interaction between email type, interruption presence, and interruption type (F(1, 314) = 20.10, $p <$ 0.001, η_p^2 = 0.06). Participants in the active condition had higher accuracy on legitimate emails regardless of interruption presence ($p_{uninterrupted} < 0.001$; $p_{interrupted} = 0.003$), while participants in the passive condition had higher accuracy on phishing emails, but only when interrupted ($p_{uninterrupted} = 0.099$; $p_{interrupted} < 0.001$).

3.2.3. Response time

To determine the effect of interruption type, email type, and interruption presence on email classification response time, a mixed ANOVA was conducted (see Fig. 5c and d). Response time was analyzed for correct trials only. Response time for interruption present trials includes the 3 s of email viewing prior to the interruption appearing and the amount of time the email was viewed after the interruption disappeared; the time the interruption was on the screen is not included in response time.

The analysis indicated no significant difference between the active and passive interruption types on email classification response time (F (1, 308) = 0.19, p = 0.667, η_p^2 = 0.0006), such that response times for participants viewing both interruption types were similar (active: M = 11.58, SD = 7.29; passive: M = 11.81, SD = 6.72). There was no main effect of email type on response time (F(1, 308) = 3.80, p = 0.052, η_p^2 = 0.01); participants took the same amount of time to respond to phishing and legitimate emails (12.59 s on average). There was a significant main effect of interruption presence on response time, however (F(1, 308) = 54.91, $p < 0.001$, $\eta_p^2 = 0.15$), such that participants took 2.22 s longer on average to classify an email when they were interrupted. There was no significant interaction between interruption presence and interruption type (F(1, 308) = 0.71, p = 0.400, η_p^2 = 0.002). Response time was similar on interrupted trials for those in the active (M = 13.31, SD = 8.11) and passive conditions (M = 13.27, SD = 9.45). A significant interaction between email type and interruption type was noted, however (F(1, 308) = 4.50, p = 0.035, η_p^2 = 0.01). Bonferroni corrected post hoc tests revealed participants in the active condition took the same amount of time to classify phishing and legitimate emails (p = 0.902), but participants in the passive condition took longer to classify legitimate emails (p = 0.005). No other interactions reached significance (all Fs < 3.14, all ps > 0.078).

3.2.4. Signal detection

As in Experiment 1, one-way repeated measures ANOVAs were conducted to identify any interruption-related differences in measures of response bias (β) and sensitivity (*d'*). A hit corresponds to a participant correctly identifying a phishing email, and a false alarm corresponds to incorrectly identifying a legitimate email as phishing. Response bias scores above 1 are considered more conservative in classifying an email as phishing, while scores below 1 are considered more liberal in classifying an email as phishing (Green and Swets, 1966/1988).

No differences between the active and passive conditions were noted for either response bias ($F(1, 314) = 3.15$, $p = 0.077$, $\eta_p^2 = 0.01$) or sensitivity ($F(1, 314) = 0.15$, $p = 0.700$, $\eta_p^2 = 0.0005$; see Fig. 6). There was however, a significant difference in response bias between interrupted and non-interrupted trials ($F(1, 314) = 14.21$, $p < 0.001$, $\eta_p^2 = 0.04$). A one-sample t-test on non-interrupted trials indicated that participants were conservative in their judgments, and thus less likely to indicate that an email was phishing ($t(315) = 5.72$, $p < 0.001$, $d = 0.32$). Alternatively, on interrupted trials, a one-sample t-test indicated that participants' beta-scores were not significantly different from 1 and therefore they were unbiased, indicating similar likelihood to classify an email as phishing or legitimate ($t(315) = 1.02$, $p = 0.310$, $d = 0.06$). Sensitivity was also significantly different between trial types ($F(1, 314) = 42.28$, $p < 0.001$, $\eta_p^2 = 0.12$); participants were better able to distinguish between phishing and legitimate emails on interrupted trials.

3.3. *Experiment 2 summary*

Experiment 2 found no group differences between active and passive interruptions on phishing classification accuracy or response time, though the two groups showed different patterns of responses. Participants in the active interruption condition had higher accuracy on legitimate emails overall, while participants in the passive interruption condition had higher accuracy on phishing emails when interrupted. On average, accuracy on phishing emails improved when participants were interrupted, but accuracy on legitimate emails did not change, though this was mainly driven by those in the passive interruption condition. In the active interruption condition, this improvement was negligible, as these participants were better at classifying legitimate emails regardless of interruption presence. In the passive condition however, participants were significantly better at classifying a phishing email when interrupted, but they were also worse at classifying legitimate emails. An overall cost of interruptions on response time was noted as well. Interrupted trials took approximately 2.2 s longer than non-interrupted trials for both groups. Across both conditions, participants were more conservative in their judgments and therefore less likely to indicate any given email was phishing on non-interrupted trials, but they were statistically unbiased on interrupted trials and therefore similarly likely to classify an email as phishing or legitimate. This was coupled with an increase in sensitivity as well. When interrupted, participants were better able to distinguish between phishing and legitimate emails.

As in Experiment 1, the increase in sensitivity and shift to unbiased responding on interrupted trials may help explain the observed increase in phishing classification accuracy when interrupted. It is possible that participants' expectations shifted on interrupted trials, such that they became more suspicious of the emails, thus causing them to respond more liberally and be more likely to classify an email as phishing. It was expected that participants in the passive interruption condition would not have to reread the emails when they were interrupted, as they would be able to maintain the email in working memory during the interruption period. This would therefore allow them to resume the main task faster and closer in proximity to where they left off than those in the active interruption condition. As there were no overall differences between the two groups however, it is possible that even those in the passive condition had to reread the emails when interrupted.

4. General discussion

Correctly and consistently identifying phishing emails is a critical task, yet it is one that most people are poor at. Previous literature has shown that individuals often miss 20–30% of phishing emails in controlled settings (Canfield, 2016; Sarno, 2020; Sheng, 2010). Reading email and identifying phishing emails are not tasks that often occur in controlled settings, however. This makes it extremely important to consider environmental factors that may impact an individual's ability to detect phishing emails. The present set of studies investigated the effects of popup task interruptions on phishing email classification. While previous research has shown the negative impact of task interruptions on other types of tasks, as well as individuals' poor overall ability to detect phishing emails, to our knowledge, this is the first study to directly assess the impact of interruptions on email classification. By requiring participants to classify emails under standard uninterrupted circumstances, as well as various interrupted conditions, the relationship between task interruptions and email classification ability could be explored.

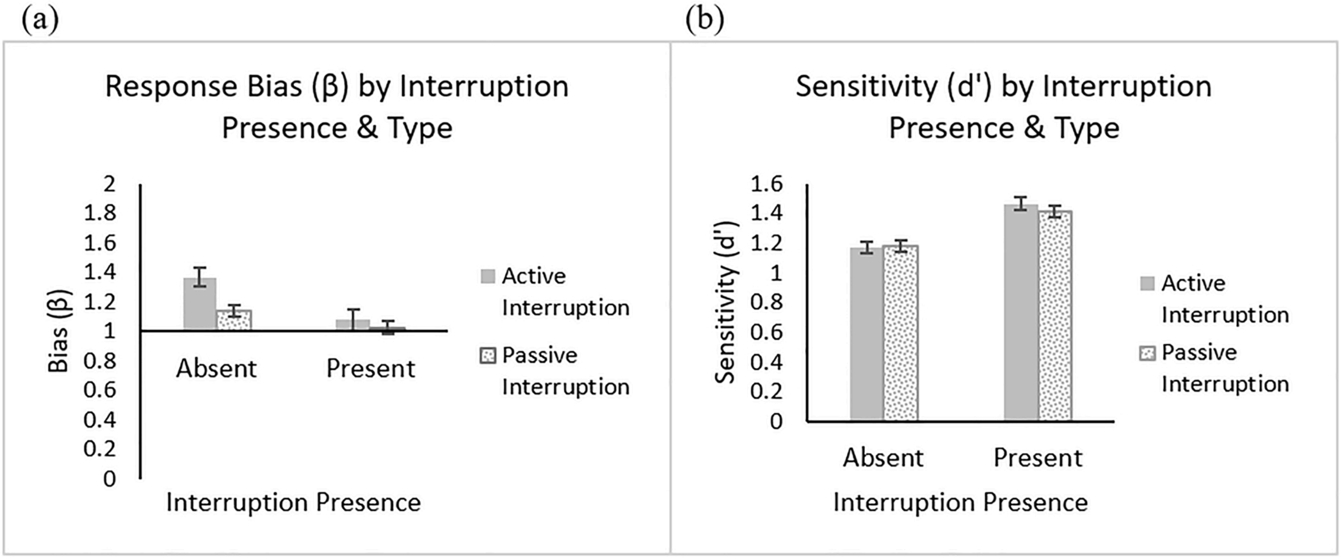

Fig. 6. Mean (a) response bias and (b) sensitivity by interruption type and presence. Error bars indicate one standard error of the mean.

4.1. Theoretical implications

Somewhat surprisingly, the findings from our studies indicated that participants were better able to classify emails as phishing when they were interrupted, though this came at the expense of longer response times. This occurred regardless of the type of interruption that participants viewed. Interrupted trials took 2.15 s longer on average than non-interrupted trials in both experiments, which equates to approximately a 20% increase in response time. This aligns with previous research indicating that interruptions increase response time on the primary task (Bailey and Konstan, 2006; Hodgetts and Jones, 2006; Williams and Drew, 2017), with one study noting up to a 27% increase (Bailey and Konstan, 2006). Overall classification accuracy was maintained on the primary task when interrupted, but phishing classification accuracy improved. Similar to Williams and Drew's (2017) finding that accuracy was maintained on a task that participants have had previous experience with, classifying an email as legitimate or phishing is a task that participants have likely had practice with, and therefore their ability to maintain accuracy when interrupted is less surprising. This maintenance is further explained by a shift in participants' response biases from conservative to unbiased when interrupted, coupled with an increase in sensitivity. Though overall accuracy on the task did not change when interrupted, more emails were correctly identified as phishing.

Adopting a working memory perspective of interruptions may provide some basis within which to contextualize the results observed in the present studies. Successfully managing interruptions while performing a task is a form of task switching. Representations of the main task must be stored in working memory to successfully return to it after engaging with an interruption, however these goal memories decay over time (Altmann and Trafton, 2002). Thus, an overwritten task goal requires more time to recall and resume (Salvucci et al., 2009). In the present experiments, it is assumed that participants maintained a representation of the email in working memory while engaging with the interruptions, as the interruptions only partially occluded the email, however this may not be accurate. The passive interruption condition in Experiment 2 was included to provide an interrupting situation in which participants did not have to complete any additional task and could continue to think about the email if they chose to do so. If participants did maintain a representation of the email in working memory during the interruption period, those in the passive interruption condition should have achieved lower classification accuracy due to retrieving their potentially degraded memory of the primary task goal over time, but shorter response times, as their representation was not hampered by competing task goals. Since there was no observable difference in performance between groups, we cannot conclusively say that the two interrupting tasks affected participant's representations of the main task in working memory any differently. This lends some credence to the memory for goals theory, as the primary task goal decayed during the interruption period, regardless of the type of interruption presented (Altmann and Trafton, 2002). Increasing the length of the interruption should therefore impact performance more severely than the current shorter interruption, though the current findings raise questions as to the nature of that impact (i.e., whether performance is helped or hindered).

Trials on which an interruption was present took approximately 2.15 s longer than non-interrupted trials in both experiments. This is close to the amount of time (3 s) that participants were able to view the email prior to being interrupted. Though an increase in response time and an increase in phishing classification accuracy were observed on interrupted trials, the interruptions themselves may not be the cause of these differences. These differences may be attributed instead to how participants reacted to, and after, the interruptions. It is possible that individuals were unable to gather enough information to form a judgment about the legitimacy of an email in those three seconds prior to the interruption occurring. Therefore, after completion of the interruption, they essentially started the initial task over (i.e., rereading the email from the beginning) to inform their decision regarding legitimacy, whether or not they maintained a representation of the email in working memory during the interruption. Repeated reading of a text has been shown to improve reading comprehension, particularly among the same passage. Rereading of one passage does not necessarily improve comprehension on a different passage, however (Therrien, 2004). If participants did reread the emails on interrupted trials, this may explain why performance improved on those specific trials. If the benefits of rereading are passage specific, any benefits gained by rereading one passage would not necessarily be observed on a different passage. This may explain the lower accuracy on non-interrupted trials when participants presumably didn't reread the emails, as any benefits gained from rereading on an interrupted trial may not be seen on an uninterrupted trial that was not read more than once. Additional studies will be required to further explore these possibilities.

4.2. Practical implications

The current studies suggest that being interrupted while reading emails, with the purpose of classifying those emails as phishing or legitimate, improves one's ability to identify phishing emails, at the expense of increased time on task. Even though accuracy improved on interrupted trials, it is important to note that participants still missed approximately 30% of phishing emails across both experiments. This is extremely problematic, as one single missed email is all it takes to cause an undesirable outcome. As noted previously, the interruptions themselves may not be what facilitated better phishing classification ability, but rather how participants reacted to the emails after being interrupted. Rereading the emails on interrupted trials may be the most effective post-interruption resumption strategy. Resuming an interrupted task goal requires backtracking, whereby people return to a previous problem state (Altmann and Trafton, 2002). Retrieving an intermediate state in a problem space requires more specific priming than returning to the original problem state. It is possible that simply seeing the email in its entirety after the interruption was not specific enough for participants to remember where they left off previously. Therefore, they returned to the original problem state: the beginning of the email. Should this be the case, cueing participants to where they left off on the email would facilitate faster and more accurate resumption after the interruption period, but would likely not increase classification accuracy. While interrupting people during an email classification task is not necessarily recommended, finding a reliable way to encourage the rereading of emails, perhaps through persistent interventions (Sarno et al., 2022), may help improve phishing classification performance.

4.3. Limitations

The current studies explore the impact of various task interruptions on an email classification task in terms of classification accuracy and overall response time. In this case, email classification was the main task that participants engaged with, but there is some question as to whether this is the primary goal of interacting with emails in the real world. It is possible that the main goal of engaging with an email is to extract information, however that goal may change to establishing legitimacy if certain characteristics are noted while reading (e.g., spelling or grammar errors, requiring an immediate response with personal information). The reason a user initially opens an email interface may influence their initial goal as well. Consequently, it is difficult to determine the initial goal of reading email and whether that goal remains constant throughout the event.

The present studies examined the difference between being interrupted with a math task compared to a blank interruption box. It is possible that the simple interrupting math task was not any more cognitively demanding than viewing a blank box. Hodgetts and Jones (2006) found that increasing the complexity of an interruption increases the time required to reinstate suspended task goals, though any interruption, even a blank screen, is enough to increase resumption time relative to uninterrupted trials. Here we analyzed a coarser measure of response time than resumption time after an interruption, so it is possible that any between-group effects were reduced to the point of insignificance.

We also observed that participants took longer to classify emails on interrupted trials, yet were better at classifying phishing emails. We speculated that this pattern might reflect participants rereading emails following interruptions, but it is unclear whether this is actually what occurred. Future studies will incorporate eye tracking measures to help elucidate this finding. Williams and Drew (2017) found that even skilled radiologists fixated on previously searched areas of radiographs after being interrupted. As reading email is a task that a majority of individuals have experience with, it is possible that being interrupted during this task will also impair memory for previously read areas of emails. Observing fixations both pre- and post-interruption will allow us to determine the extent to which people reread emails when interrupted.

4.4. Conclusion

Overall, the present studies have provided some evidence that limited popup interruptions improve phishing email classification. While this may be encouraging news, it is important to remember that participants were nowhere near perfect at this task. When it comes to the correct classification of phishing emails, one single missed phishing email is enough to cause disaster. The goal is to elevate performance to where individuals can correctly classify every phishing email every time, and our findings may suggest a candidate pathway toward that goal. While using interruptions as a means of improving performance may seem counterintuitive, the effects observed here suggest that improved performance may be less related to the interruptions themselves, and more to the robustness of the processing of the email to be classified. More specifically, the presence of interrupting tasks in our studies, regardless of their nature, may have encouraged better reading comprehension. In this light, the improved performance we observed when interruptions were present seems a little less counterintuitive. Classifying an email effectively requires that the content of the email be well read and understood, and the rereading that task interruptions seem to have elicited in our studies likely supports those requirements. Future studies are needed to critically evaluate the robustness of this finding and its potential to be utilized as a training intervention for phishing awareness and classification.

CRediT authorship contribution statement

Elisabeth J.D. Slifkin: Conceptualization, Methodology, Software, Investigation, Formal analysis, Data curation, Visualization, Writing – original draft, Project administration. **Mark B. Neider:** Conceptualization, Methodology, Writing – review & editing, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

Data availability

The authors do not have permission to share data.

References

Adamczyk, P.D., Bailey, B.P., 2004. If not now, when? The effects of interruption at different moments within task execution. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04, pp. 271–278. https://doi.org/10.1145/985692.985727.

Altmann, E.M., Trafton, J.G., 2002. Memory for goals: an activation-based model. Cogn. Sci. 26 (1), 39–83. https://doi.org/10.1207/s15516709cog2601_2.

Altmann, E.M., Trafton, J.G., Hambrick, D.Z., 2014. Momentary interruptions can derail the train of thought. J. Exp. Psychol. General 143 (1), 215. https://doi.org/10.1037/a0030986.

Bailey, B.P., Konstan, J.A., 2006. On the need for attention-aware systems: measuring effects of interruption on task performance, error rate, and affective state. Comput. Hum. Behav. 22 (4), 685–708. https://doi.org/10.1016/j.chb.2005.12.009.

Bergholz, A., De Beer, J., Glahn, S., Moens, M.F., Paaß, G., Strobel, S., 2010. New filtering approaches for phishing email. J. Comput. Secur. 18 (1), 7–35. https://doi.org/10.3233/JCS-2010-0371.

Boehm-Davis, D.A., Remington, R., 2009. Reducing the disruptive effects of interruption: a cognitive framework for analysing the costs and benefits of intervention strategies. Accident Anal. Prev. 41 (5), 1124–1129. https://doi.org/10.1016/j.aap.2009.06.029.

Canfield, C.I., Fischhoff, B., Davis, A., 2016. Quantifying phishing susceptibility for detection and behavior decisions. Hum. Factors 58 (8), 1158–1172. https://doi.org/10.1177/0018720816665025.

Ceci, L. 2022. Email usage in the United States—statistics & facts. Statista. https://www.statista.com/topics/4295/e-mail-usage-in-the-united-states/.

Couffe, C., Michael, G.A., 2017. Failures due to interruptions or distractions: a review and a new framework. Am. J. Psychol. 130 (2), 163–181. https://doi.org/10.5406/amerjpsyc.130.2.0163.

Drake, C.E., Oliver, J.J., Koontz, E.J., 2004. Anatomy of a Phishing Email. CEAS.

Foroughi, C.K., Werner, N.E., Barragán, D., Boehm-Davis, D.A., 2015. Interruptions disrupt reading comprehension. J. Exp. Psychol. General 144 (3), 704. https://doi.org/10.1037/xge0000074.

Golladay, K., Holtfreter, K., 2017. The consequences of identity theft victimization: an examination of emotional and physical health outcomes. Vict. Offenders 12 (5), 741–760. https://doi.org/10.1080/15564886.2016.1177766.

Gangavarapu, T., Jaidhar, C.D., Chanduka, B., 2020. Applicability of machine learning in spam and phishing email filtering: review and approaches. Artif. Intell. Rev. 1–63. https://doi.org/10.1007/s10462-020-09814-9.

Gillie, T., Broadbent, D., 1989. What makes interruptions disruptive? A study of length, similarity, and complexity. Psychol. Res. 50 (4), 243–250. https://doi.org/10.1007/BF00309260.

Gorham, M., 2020. 2019 Internet Crime Report (FBI Internet Crime Report). United States, Federal Bureau of Investigation, Internet Crime Complaint Center.

Green, D.M., Swets, J.A., 1966/1988. Signal Detection Theory and Psychophysics, 1. Wiley, New York.

Grimes, G.A., Hough, M.G., Signorella, M.L., 2007. Email end users and spam: relations of gender and age group to attitudes and actions. Comput. Hum. Behav. 23 (1), 318–332. https://doi.org/10.1016/j.chb.2004.10.015.

Hodgetts, H.M., Jones, D.M., 2006. Interruption of the tower of London task: support for a goal-activation approach. J. Exp. Psychol. General 135 (1), 103–115. https://doi.org/10.1037/0096-3445.135.1.103.

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. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 57, pp. 1012–1016. https://doi.org/10.1177/1541931213571226.

Identity Theft Resource Center, 2021. ITRC consumer aftermath report: how identity crimes impact victims, their families, friends, and workplaces. [PowerPoint slides]. https://www.idtheftcenter.org/wp-content/uploads/2021/09/ITRC_2021_Consumer_Aftermath_Report.pdf.

Iqbal, S.T., Horvitz, E., 2007. Disruption and recovery of computing tasks: field study, analysis, and directions. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 677–686. https://doi.org/10.1145/1240624.1240730.

Kumaraguru, P., Rhee, Y., Acquisti, A., Cranor, L.F., Hong, J., Nunge, E., 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, pp. 905–914. https://doi.org/10.1145/1240624.1240760.

Leroy, S., Glomb, T.M., 2018. Tasks interrupted: how anticipating time pressure on resumption of an interrupted task causes attention residue and low performance on interrupting tasks and how a "ready-to-resume" plan mitigates the effects. Organ. Sci. 29 (3), 380–397. https://doi.org/10.1287/orsc.2017.1184.

Li, W., Lee, J., Purl, J., Greitzer, F., Yousefi, B., Laskey, K., 2020. Experimental investigation of demographic factors related to phishing susceptibility. In: Proceedings of the 53rd Hawaii International Conference on System Sciences, pp. 2240–2249. https://doi.org/10.24251/HICSS.2020.274.

Li, Y., Yazdanmehr, A., Wang, J., Rao, H.R., 2019. Responding to identity theft: a victimization perspective. Decis. Support Syst. 121, 13–24. https://doi.org/10.1016/j.dss.2019.04.002.

Mark, G., Iqbal, S.T., Czerwinski, M., Johns, P., Sano, A., Lutchyn, Y., 2016. Email duration, batching and self-interruption: Patterns of email use on productivity and stress. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 1717–1728. https://doi.org/10.1145/2858036.2858262.

Mayhorn, C.B., Nyeste, P.G., 2012. Training users to counteract phishing. Work 41 (Supplement 1), 3549–3552. https://doi.org/10.3233/WOR-2012-1054-3549.

Morey, R.D., 2008. Confidence intervals from normalized data: a correction to Cousineau, 2005 Tutor. Quant. Methods Psychol. 4 (2), 61–64. https://doi.org/10.20982/tqmp.04.2.p061.

Nyeste, P.G., Mayhorn, C.B., 2010. Training users to counteract phishing. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 54, pp. 1956–1960. https://doi.org/10.1177/154193121005402311.

O'Conaill, B., Frohlich, D., 1995. Timespace in the workplace: dealing with interruptions. In: Proceedings of the Conference on Human Factors in Computing Systems (CHI '95), pp. 262–263.

Oliveira, D., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., Muradoglu, M., Weir, D., Soliman, A., Lin, T., Ebner, N., 2017. Dissecting spear phishing emails for older vs young adults: on the interplay of weapons of influence and life domains in predicting susceptibility to phishing. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, pp. 6412–6424. https://doi.org/10.1145/3025453.3025831.

Oulasvirta, A., Saariluoma, P., 2006. Surviving task interruptions: investigating the implications of long-term working memory theory. Int. J. Hum. Comput. Stud. 64 (10), 941–961. https://doi.org/10.1016/j.ijhcs.2006.04.006.

Patel, P., Sarno, D.M., Lewis, J.E., Shoss, M., Neider, M.B., Bohil, C.J., 2019. Perceptual representation of spam and phishing emails. Appl. Cogn. Psychol. 33 (6), 1296–1304. https://doi.org/10.1002/acp.3594.

Peirce, J., Gray, J.R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J.K., 2019. PsychoPy2: experiments in behavior made easy. Behav. Res. 51, 195–203. https://doi.org/10.3758/s13428-018-01193-y, 2019.

Pew Research Center, 2021, April 7. Mobile fact sheet. Pew Research Center. https://www.pewresearch.org/internet/fact-sheet/mobile/.

Puranik, H., Koopman, J., Vough, H.C., 2020. Pardon the interruption: an integrative review and future research agenda for research on work interruptions. J. Manag. 46 (6), 806–842. https://doi.org/10.1177/0149206319887428.

Ratwani, R.M., Andrews, A.E., McCurry, M., Trafton, J.G., Peterson, M.S., 2007. Using peripheral processing and spatial memory to facilitate task resumption. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 51, pp. 244–248. https://doi.org/10.1177/154193120705100421.

Ratwani, R.M., Trafton, J.G., 2008. Spatial memory guides task resumption. Visual Cogn. 16 (8), 1001–1010. https://doi.org/10.1080/13506280802025791.

Salvucci, D.D., Taatgen, N.A., Borst, J.P., 2009. Toward a unified theory of the multitasking continuum: from concurrent performance to task switching, interruption, and resumption. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 1819–1828. https://doi.org/10.1145/1518701.1518981.

Sarno, D.M., Lewis, J.E., Bohil, C.J., Shoss, M.K., Neider, M.B., 2017. Who are phishers luring?: A demographic analysis of those susceptible to fake emails. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 61, pp. 1735–1739. https://doi.org/10.1177/1541931213601915.

Sarno, D.M., Lewis, J.E., Bohil, C.J., Neider, M.B., 2020. Which phish is on the hook? Phishing vulnerability for older versus younger adults. Hum. Factors 62 (5), 704–717. https://doi.org/10.1177/0018720819855570.

Sarno, D.M., McPherson, R., Neider, M.B., 2022. Is the key to phishing training persistence?: Developing a novel persistent intervention. J. Exp. Psychol. Appl. 28 (1), 85–99. https://doi.org/10.1037/xap0000410.

Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L.F., Downs, J., 2010. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 373–382. https://doi.org/10.1145/1753326.1753383.

Speier, C., Valacich, J.S., Vessey, I., 1999. The influence of task interruption on individual decision making: an information overload perspective. Decision Sci. 30 (2), 337–360. https://doi.org/10.1111/j.1540-5915.1999.tb01613.x.

Therrien, W.J., 2004. Fluency and comprehension gains as a result of repeated reading: a meta-analysis. Remedial Special Educ. 25 (4), 252–261. https://doi.org/10.1177/07419325040250040801.

Trafton, J.G., Altmann, E.M., Brock, D.P., Mintz, F.E., 2003. Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. Int. J. Hum. Comput. Stud. 58 (5), 583–603. https://doi.org/10.1016/S1071-5819(03)00023-5.

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. Decis. Support Syst. 51 (3), 576–586. https://doi.org/10.1016/j.dss.2011.03.002.

Vishwanath, A., Harrison, B., Ng, Y.J., 2018. Suspicion, cognition, and automaticity model of phishing susceptibility. Commun. Res. 45 (8), 1146–1166. https://doi.org/10.1016/j.dss.2011.03.002.

Williams, L.H., Drew, T., 2017. Distraction in diagnostic radiology: how is search through volumetric medical images affected by interruptions? Cogn. Res. Princ. Implic. 2 (1), 1–11. https://doi.org/10.1186/s41235-017-0050-y.

Williams, S.E., Sarno, D.M., Lewis, J.E., Shoss, M.K., Neider, M.B., Bohil, C.J., 2019. The psychological interaction of spam email features. Ergonomics 62 (8), 983–994. https://doi.org/10.1080/00140139.2019.1614681.