

User Context: An Explanatory Variable in Phishing Susceptibility

Kristen K. Greene
National Institute of Standards
and Technology
kristen.greene@nist.gov

Michelle P. Steves
National Institute of Standards
and Technology
michelle.steves@nist.gov

Mary F. Theofanos
National Institute of Standards
and Technology
mary.theofanos@nist.gov

Jennifer Kostick
National Institute of Standards
and Technology
jennifer.kostick@nist.gov

Abstract—Extensive research has been performed to examine the effectiveness of phishing defenses, but much of this research was performed in laboratory settings. In contrast, this work presents 4.5 years of workplace-situated, embedded phishing email training exercise data, focusing on the last three phishing exercises with participant feedback. The sample was an operating unit consisting of approximately 70 staff members within a U.S. government research institution. A multiple methods assessment approach revealed that the individual’s work context is the lens through which email cues are interpreted. Not only do clickers and non-clickers attend to different cues, they interpret the same cues differently depending on the alignment of the user’s work context and the premise of the phishing email. Clickers were concerned over consequences arising from *not* clicking, such as failing to be responsive. In contrast, non-clickers were concerned with consequences from clicking, such as downloading malware. This finding firmly identifies the alignment of user context and the phishing attack premise as a significant explanatory factor in phishing susceptibility. We present additional findings that have actionable operational security implications. The long-term, embedded and ecologically valid conditions surrounding these phishing exercises provided the crucial elements necessary for these findings to surface and be confirmed.

Keywords—decision-making, embedded phishing awareness training, user-centered approach, survey instrument, long-term assessment, operational data, trial deployment, network security, security defenses

I. INTRODUCTION

The problem of phishing is not solved. It is an escalating cyber threat facing organizations of all types and sizes, including industry, academia, and government [1], [11], [20], [22]. Often using email, phishing is an attempt by a malicious actor posing as trustworthy to install malware or steal sensitive information for financial gain. The nature of phishing itself has changed, moving far beyond “traditional” phishing for usernames, passwords, and credit card numbers via fraudulent websites, and into more sophisticated cybercrime attacks that mount advanced persistent threats against organizations and steal individuals’ financial identities with devastating consequences for both users and organizations. The practice of phishing has turned a pervasive means of communication—email—into a dangerous threat channel. Symantec reports that malicious emails were the weapon of choice for bad actors, ranging from state-sponsored espionage groups to mass-mailing ransomware gangs, and that one in 131 emails sent during 2016 was malicious [22].

Advanced threats via ransomware increased 167 times from four million attempts in 2015 to 638 million attempts in 2016, mostly through phishing campaigns [20]. Email’s popularity with attackers is driven by several factors. It is a proven attack vector. It does not rely on system vulnerabilities, but on human deception. Routine business processes, such as correspondence about delivery notifications and invoices, provide camouflage for these malicious emails and were the favored guise for spreading ransomware in 2016 [22]. In the information security domain, the use of deception to manipulate individuals for fraudulent purposes is referred to as social engineering.

To help combat the phishing threat, many organizations utilize some type of phishing awareness training to make employees and students more aware of phishing threats and consequences, e.g., Stanford University’s Phishing Awareness Service [21]. These embedded phishing awareness training systems use software to send simulated phishing emails to users’ regular email accounts. By “phishing” users in their normal computing environments, these emails are intended to train people to recognize and avoid falling victim to phishing attacks in their work (or school) setting. Emails are designed to emulate real-world threats currently facing organizations, providing a realistic experience in a safe, controlled way so recipients can become familiar with the types of tactics used in real phishing attacks. Embedded training schemes that combine training people in their normal work environments with immediate feedback produce more lasting change to behaviors and attitudes [14]. It also provides the means to capture click decisions in an operational environment.

The goal of this work was to better understand why users click and do not click on links or open attachments in phishing training emails that were part of embedded phishing awareness exercises at the National Institute of Standards and Technology (NIST), a U.S. government research institution. From mid-2012 through 2015, the institute’s Information Technology Security and Networking Division (ITSND) facilitated 12 operationally-situated exercises using a commercially-available phishing awareness training system. All exercises were conducted in one particular operating unit (OU) at the institute. Although staff knew their OU was participating in these exercises, they were not announced and were conducted at irregular intervals to avoid priming effects. Unfortunately, click rates were variable across the initial 3.5 years of exercises, making the training effect difficult to characterize. In 2016, human factors researchers partnered with the institute’s ITSND to better understand the variability in the operational click-rate data. Over the course of 2016, another three exercises were conducted exactly as before with the following exception: each exercise also had an accompanying post-exercise survey to better understand why

users were clicking or not clicking. The long-term data from the 2016 exercises alone represents real-world phishing data with breadth and depth unlike any reported in the literature to-date. For all exercises in the 4.5-year span (2012 through 2016), the phishing training emails modeled real-world phishing campaigns, and participating staff were in their normal work environments with their regular work loads, providing ecological validity. This paper presents the novel finding that the alignment of user context and the phishing message premise is a significant explanatory factor in phishing susceptibility, impacting depth of processing and concern over consequences. We believe that the rare opportunity to collect data surrounding phishing click decisions in the workplace coupled with ecologically valid conditions over time provided the crucial elements for these findings to surface. Because of the scarcity of operationally-situated data, we also report on findings we believe to be actionable now in operational settings.

II. BACKGROUND

A. Background research

Technological and human-centered approaches are used to combat email phishing. Technology-based solutions generally focus on reducing software vulnerabilities, for example, maintaining software currency and identifying malicious websites and phishing emails based on their characteristics. This identification centers around using server-side filtering and classifiers and client-side filtering tools [2]. The server-side mechanisms strive to remove malicious messages and website links before the user sees them, while the client-side tools often attempt to aid user decision-making [9], [11]. Much of the filter algorithm development and classifier training is done offline, prone to error, and reactive in nature [2].

Human-centered approaches attempt to bridge the gap left by reactive technological solutions [9]. Research in these approaches often falls into one of three categories: educational awareness training to identify phishing messages, new user interface mechanisms and designs coupled with client-side filtering intended to aid users' click decisions, and research that considers psychological factors in decision-making with respect to phishing [18].

While there are many studies that have explored the ability of email users to recognize phishing cues, most of these have been conducted in laboratory settings where users are not faced with click decisions in real-world settings under their normal workloads and time pressures, and without laboratory priming effects. In contrast to laboratory-based studies, our user-centered phishing assessment situates participants in the intended use setting: the employee workplace.

There are only a few studies in the literature using embedded, simulated phishing in the user's normal computing environment during their normal work day: [4], [8], [16], and [17]. These studies were set in the real world using operationally-situated study settings, but were focused on investigating training materials and click rates, not on participant click-decision factors. The studies described in [4] and [16] primarily looked at the efficacy of embedded awareness training. Those reported in [8] and [17] focused on the startling number of users who clicked in their respective operational environments. Click decision exploration was not included in [16] and [17], and only speculated about in [8].

The workplace-situated study reported in [4] is the most similar study to our work of those cited, although the focus of that study was to examine the effect of training materials on click decisions rather than other factors surrounding these decisions. Caputo et al. describe three phishing training exercises over the course of eight months. Only after the third exercise was completed were interviews conducted with 27 participants. Although cursory thematic data were reported, details about click decisions were not given, with the caveat that almost all interviewees who clicked most often recalled the third trial only, while most non-clickers did not recall any of the training phish due to the length of time that elapsed between the initial phish and the interview. In contrast, our work centers on investigating factors surrounding participant click decisions rather than training material effectiveness.

There are many reasons why there are so few studies set in the real world. [4], [8], and [16], among others, note challenges in conducting real-world phishing studies that can provide more ecological validity and richer data than those administered in laboratory settings. For example, coordinating with operational staff to get training phish through corporate firewalls and filters are hurdles that must be overcome. Maintaining participant privacy and avoiding participant cross-contamination, such as warning other participants [15], are note-worthy challenges. The expense of examining training retention for longer than a 90-day period [4] and obtaining stakeholder buy-in [8] are also mentioned. Given these significant challenges, why attempt to study phishing click decisions *in the workplace*? The answer centers on *context of use*—how, where, and under what circumstances email users are making click decisions. Indeed, Wang, et al. note the enormous contribution data from real phishing victims would provide, if it were available [27]. The data presented here provide rare insights into click decisions in the workplace.

B. Project background

This project was started in 2012 by the institute's ITSND as a long-term trial deployment of an embedded phishing awareness training effort. The trial deployment was intended as a multi-year effort and used a commercially-available system to help develop and deliver phish messages and training, as well as track click rates. The same system was used throughout the entire 4.5-year period. For all exercises, the targeted population within the institute was one operating unit having approximately 70 staff members. The awareness training provided by these exercises augmented the IT security awareness training the entire institute's workforce received annually. OU staff were aware their unit was participating in the trial. However, exercises were unannounced and deployed at irregular intervals to avoid priming effects.

All exercises were conducted by the OU's Information Technology Security Officer (ITSO). The same person held the position during the entire 4.5-year period. The ITSO selected the phishing message and its premise from templates provided by the training system that mimicked current real-world threats. The ITSO used some messages without modification so response rates could be compared with other organizations using the same training system and message, a form of benchmarking. Some messages were tailored to align with business and communication practices within the organization or were personalized, in other words, they were spear phish [25], [27].

The ITSO also developed the training given to those who clicked. Further, the ITSO set the timing of each exercise and coordinated with the IT security team to allow the phish through the institute’s firewalls and filtering mechanisms.

Fig. 1 shows the click rates for the exercises conducted in 2012 through 2016. For the first 12 exercises, those without survey feedback, click rates ranged from 1.6 % to 49.3 %, having a *Mean* = 17.3 %, *Median* = 11.9 %, and *SD* = 15.2. The social engineering premise for each exercise was varied, except for the three ‘Package’ exercises that used the same message mimicking a package delivery notice. This benchmark phish was not tailored to the organization. Click rates for the ‘Package’ exercises were 28.6 %, 12.2 %, and 7.7 %. Interestingly, despite being seen three times, the click rate for this phish exercise did not go to zero, although it did decline. All premises used imitations of normal discourse [3] and were based on then-current social engineering scams. The sample was the OU; the sole inclusion criterion was being assigned to the OU. The only OU staff member excluded was the OU’s ITSO, who conducted the exercises. The OU had an annualized separation rate of 8.7 % during the entire 4.5-year period. Individuals were intentionally not tracked across exercises to protect employee privacy.

Although click rates declined on average during the first 3.5 years of training, the trend did not approach zero. Given the puzzling variability in click rates, the click-rate data alone from 12 exercises over the 3.5-year period were insufficient to characterize the training effect. Therefore, in early 2016, human factors researchers employed by the institute were asked to help investigate and explain the variability in click-rate data from the trial deployment. A review of click rates and scenario types did not yield a satisfactory explanation, only more questions. Quantitative click-rate data only told us *how many* people clicked, but not *why* they clicked. We realized we needed click-decision information directly from participants to help explain why they were clicking and not clicking in their own words. Garfinkel and Lipford expressed it this way, “In order to train users to avoid phishing attacks it is necessary to first understand why people are falling for them” [9].

III. METHOD

To better interpret the institution’s prior 3.5 years of operational click-rate data—and understand click rates for future planned exercises—we decided on a multiple methods assessment approach [24]. We knew that numbers alone did not

tell the whole story and that purely quantitative click-rate data could not answer our assessment question: *Why are email users clicking or not clicking on phishing links and attachments?* Further, we needed to understand why people were or were not clicking to inform the organization’s trial deployment question: *Why are click rates so variable?* We decided to use a survey instrument to probe click decisions by obtaining participant feedback following each of the next three training exercises in 2016, represented in Fig. 1 as triangle-shaped data points. The survey method allowed for immediate administration of the instrument following a click decision. Additionally, an anonymous, online survey was preferable in the workplace to provide the freedom for more honest responses compared to in-person methods. We followed the appropriate human subjects approval process for our institution.

A. Approach

We developed, tested, and fielded a web-based survey instrument for three new phishing awareness training exercises in 2016, while meeting regularly with our ITSO partner. Our multiple methods assessment used a survey instrument having open- and closed-ended responses to gather qualitative and quantitative data to complement our click-rate data. Our coding and analysis approach probed similarities and differences in survey response data between and among clickers and non-clickers.

When developing and conducting phishing exercises during 2016, the ITSO followed the same protocol as previously described for the 2012 to 2015 exercises, with the following exception: the ITSO sent a survey invitation email to each employee after they clicked or did not click the link or attachment in the phishing training email. The phishing training software tracked whether participants clicked or not. For those who clicked, the ITSO immediately sent a survey invitation email. For those who did not click, the ITSO waited a week before sending the survey invitation email (to allow them sufficient time to have seen the phishing exercise email). The survey invitation email included the names and contact information of the researchers conducting the assessment, and the appropriate link depending on the employee’s click decision. Although operational security data on click rates were identifiable, survey data were confidential. Intentionally, there was no linkage between individuals and survey responses to help preserve privacy in the workplace.

B. Environment and participants

It is always important to understand the environment in which a study is conducted, but this is especially critical for a long-term operational assessment conducted in situ. The institute where this assessment took place is highly security-conscious, with yearly mandatory IT security awareness training for all staff. The OU that participated in the deployment benefited from an involved ITSO. For instance, the ITSO proactively sent emails advising people of current security scams, developed phishing awareness training that was tailored

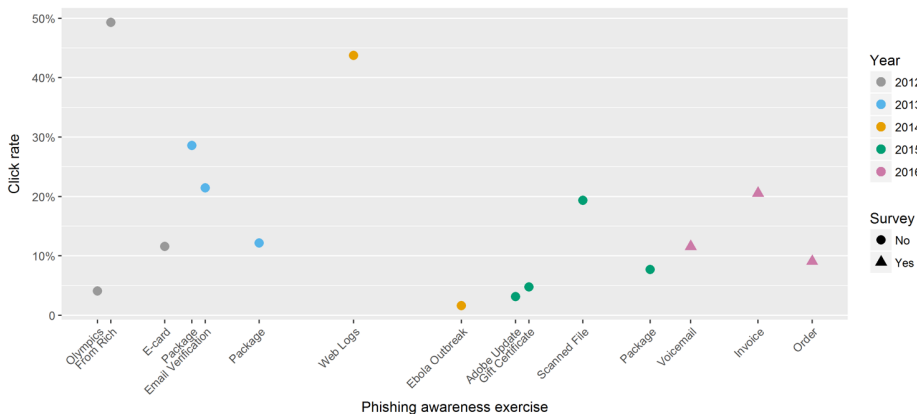


Fig 1: Phishing Awareness Exercise Click Rates, 2012 through 2016

specifically to the OU, reminded staff to forward suspicious emails to the security team, and generally encouraged people to be actively engaged in security. Our ITSO partner was well-known within the OU and staff knew of the phishing emails. The fact that staff were aware they were participating in a long-term phishing awareness effort differed from the studies reported in [4], [8], and likely [17], where those assessed did not know they would be phished by their respective organizations.

Participants in our assessment sample were staff at a U.S. government research institution, specifically those individuals comprising the OU previously described. The OU had good variability in age, education, supervisory experience, and time spent at the research institution. The OU was 67 % male and 33 % female, with 23 % of individuals in supervisory positions and 77 % in non-supervisory positions. Ages ranged from 25 to 66 years, $Mean = 48.7$, $SD = 9.9$. Education ranged from high school to doctorate degrees. Time working at the research institution ranged from less than one year to 39 years. The OU demographics were similar to those of the larger institution, which has approximately 3000 staff. N , the number of people who participated in a given phishing exercise, ranged from 66 to 73 across the three training exercises, as shown in Table I of the Results section.

C. Instrument development

In consultation with our ITSO partner, we developed an online survey instrument for the March 2016 phishing awareness training exercise. We solicited feedback on our survey from three domain experts; these were Human-Computer Interaction (HCI) researchers with significant usable security domain expertise. Based on the resulting survey data from the March exercise, and to customize the survey for upcoming exercises, we made minor refinements to the survey instrument template for use in the August and December exercises. Slight refinements were made to refer to an attachment versus a link (as appropriate given the nature of each exercise) and to include the screen image capture of the phishing email in question. We also added specific close-ended questions to further probe findings from open-ended responses in the first survey. E.g., we added a four-point Likert scale question on confidence in institutional computer security measures (with an open-ended follow-up question), and a yes/no/not sure question on perceived behavior change (also with an open-ended follow-up). To further validate the refined survey instrument [13], the instrument was reviewed by two survey experts and two additional domain experts. We also performed two formal cognitive walkthroughs with pseudo-participants (persons representative of the participant sample, but who were not part of the sample).

Importantly, the structure and content of all three surveys was consistent and nearly identical, with the aforementioned customizations made to the survey template based on the nature of the phishing email (link versus attachment). This is especially critical when considering the open-ended qualitative data from participants' impressions of the email and click decision explanations, which form the bulk of our analyses and findings. These were consistent across exercises. In order to better elicit participants' initial impressions without bias, these open-ended questions started the survey template, with close-ended questions following thereafter. Clickers received a survey version that said, "did click" and non-clickers received a version

that said, "did not click" when referring to the open-ended questions about initial impressions and click/non-click decisions. That was the only wording difference between clickers and non-clickers; all other questions were identical.

Surveys were conducted via commercially-available web-based software. The survey template consisted of a mixture of open-response (3 questions) and close-ended questions (17 questions, 7 of which had open-ended follow-ups), for a total of 20 questions and an estimated response time of only 9 minutes. We felt it was important to have a short survey since participants were in the work setting and we wished to minimize time taken away from their regular tasks. Survey questions addressed initial impressions of the phishing training email, click decision explanations, device used, concern over possible consequences for clicking/not clicking, hurrying, curiosity, suspicion, confidence in the institution's computer security measures, number of emails sent and received daily, demographics, impressions of phishing awareness training and behavior change, and a final question asking for any additional comments on the survey or phishing exercises. It is important to reiterate that questions on initial impressions and click decision explanations were open-ended and consistent across exercises. Additionally, as should be clear in the Results section, there were open-ended follow-ups to questions on confidence in institutional computer security measures, consideration of consequences, and behavior change. We encourage other researchers to use our survey instrument template (Appendix Table 1) for future operational phishing research.

D. Training exercise details

Here we describe the three phishing awareness training exercises with corresponding surveys conducted during 2016. Each exercise used a different premise based on a real-world phishing threat current at that time. Two exercises were link attack scenarios, designed to mimic general malware distribution attacks. One exercise was an attachment attack scenario designed specifically to mimic the Locky ransomware attack. Each email contained cues, such as misspellings or a missing salutation, that a discerning recipient could have used to correctly identify the email as a phish. None of the three emails asked for personal information such as usernames or passwords, often a suspicion trigger [7], as such traditional phishing attacks were not the most prevalent threats at that time.

1) First exercise: new voicemail (March)

In the first exercise, the phishing training email appeared to be a system-generated message from the fictitious CorpVM <mailto:corpvm@webaccess-alert.com>. The subject line was "You have a new voicemail." The greeting was personalized with the recipient's first and last name. The body referred to an undelivered voicemail with these instructions: "To listen to this message, please click here. You must have speakers enabled to listen to the message." The body also included a reference number for the message, length of transmission, receiving machine ID, thank-you, and smaller footer text that said, "This is a system-generated message from a send-only address. Please do not reply to this email."

2) Second exercise: unpaid invoice (August)

In the second exercise, the phishing training email appeared to be from a fictitious federal employee of the same institution

as the recipients: Preston, Jill (Fed) <jill.preston@nist.gov>. The subject line was “Unpaid invoice #4806.” The greeting was personalized with the recipient’s first and last name. The body of the email said, “Please see the attached invoice (.doc) and remit payment according to the terms listed at the bottom of the invoice. Let us know if you have any questions. We greatly appreciate your prompt attention to this matter!” The email closed with the name of the fictitious federal employee (“Jill Preston”), but without any additional contact information. The attachment was labeled as a .zip, although the email body text referenced the filename “invoice_S-37644806.doc”—an extension mismatch.

3) Third exercise: order confirmation (December)

In the third exercise, the phishing training email appeared to be from the fictitious “Order Confirmation auto-confirm@discontcomputers.com” [sic] with this subject line: “Your order has been processed.” There was no greeting, either personalized or generic. The body said, “Thank you for ordering with us. Your order has been processed. We’ll send a confirmation e-mail when your item ships.” The body also included an image of holiday packages—this email was sent in December before the holidays, order details, an order number, estimated delivery date, subtotal, tax, and order total. There was a “Manage order” button that users could click and closing text that said, “Thank you for your order. We hope you return soon for more amazing deals.”

E. Data analysis

Throughout the data analysis process, we focused on answering our operational assessment question: *Why are email users clicking or not clicking on phishing links and attachments?* With our multiple methods approach, we used qualitative and quantitative survey data from participants, quantitative click-rate data captured automatically by the phishing training software, and observations by the phishing awareness training conductor—our ITSO partner.

As our open-ended survey questions were few, short, and bounded, and the response space was small, we chose to use a single coder rather than engage in group coding. Using an iterative analysis process [19], the same team member coded all open-ended survey data then reported back to the other research team members for analytic group discussions. More specifically, the coder first read through the survey data multiple times to become familiar with the content before beginning the process of coding, then prepared spreadsheets to help examine response similarities and differences between clickers and non-clickers for each question with respect to emerging themes. All research team members met regularly to discuss the ways in which codes led to particular themes, and to identify relationships amongst codes.

The coder also computed descriptive statistics for the quantitative data, comprised of click-rate data and close-ended survey data. These quantitative analyses were also shared with the larger research team and examined in conjunction with the qualitative coding during analytic group discussions. The entire team examined the qualitative and quantitative analyses for each of the three phishing exercises and compared findings across exercises, focusing on similarities and differences in clickers’ and non-clickers’ responses among the phishing exercises. These comparisons across exercises showed good triangulation

within and across clickers and non-clickers. Finally, we performed member-checking by presenting our analyses—both qualitative and quantitative—to our ITSO partner, who provided us with contextual information about each phishing exercise.

We chose to focus the following Results section on our qualitative results and themes, as these were novel findings surrounding user context over nearly a year in a true operational setting. As previously described, our open-ended survey questions asked participants in their own words to describe their initial impressions of the phishing email and to explain why they chose to click/not click on the phishing link or attachment.

IV. RESULTS

For the three 2016 phishing exercises, Table I presents click rates—the number of recipients who clicked on the link or opened the attachment—and survey response rates. The unpaid invoice

TABLE I. PHISHING EXERCISE STATISTICS

Exercise	<i>n</i>	Phishing click rate	Survey response rate, clickers	Survey response rate, non-clickers
New voicemail	69	11.6 % (8/69)	100 % (8/8)	21.3 % (13/61)
Unpaid invoice	73	20.5 % (15/73)	66.7 % (10/15)	25.9 % (15/58)
Order confirmation	66	9.1 % (6/66)	66.7 % (4/6)	50 % (30/60)

scenario had by far the highest click rate, approximately twice that of the other two scenarios. In subsequent sections, we present a summary of findings for each exercise, followed by primary themes that ran across exercises. Illustrative supporting quotes with reference codes are intermixed throughout. Each participant was assigned a reference code, denoted as ###C or ###NC for clickers and non-clickers, respectively. Numbers in the 100 series refer to the first exercise, the 200 series refers to the second exercise, and the 300 series refers to the third exercise. We focus heavily on results from the second exercise for several important reasons: 1) this phishing email mimicked the real-world Locky ransomware prevalent at that time, 2) it had the highest click rate of the three exercises, and 3) we therefore have the most survey data from clickers for that exercise.

A. Phishing premises and cues

For both clickers and non-clickers, there seemed to be an accumulation of cues that contributed to their click/non-click decisions in each exercise, with an individual’s work context being the lens through which all other cues were viewed and interpreted. By user context, we mean the unique combination of an individual’s work setting, responsibilities, tasks, and recent real-world events in their life: how a person experiences reality in the workplace. Most importantly, the alignment between a user’s work context and the general premise of a phishing email determined whether they found the email initially believable or suspicious, which in turn influenced the cues they attended to.

1) First exercise: new voicemail (March)

Clickers perceived the email as looking legitimate and professional, and aligned with external events, such as having

missed a call earlier. They were unsuspecting because they were not asked for personal information, and they found the email believable because unfamiliar emails at work are common for them. For example, when explaining their click decisions, respondent 103C stated, “My phone rang earlier but I wasn’t able to pick up in time,” and 106C noted, “Recent talks about changing to a VOIP system.” (VOIP stands for Voice over Internet Protocol.) 101C said, “The email looks professional,” and 105C explained, “It looked legitimate.” 104C elaborated, “The unfamiliar email is common at work, and generally not a problem. Did not trigger anything in my brain that would indicate that it was harmful. Did not ask me to give personal info like social security number etc.”

In contrast, non-clickers were suspicious of the email’s appearance, often describing it as “strange,” “spammy,” and “unusual.” They also reported suspicion-raising cues in the *from* address or company name; misalignment with expectations or external events (no voicemail indicator on phone); and the unfamiliar message, never having received an email about a voicemail before. For example, 102NC said, “this had a very ‘spammy’ feel to it, especially since it was a service I’ve never heard of.” 105NC noted, “it did not look like an official government email.” 109NC explained, “It was from a ‘.com’ email address instead of ‘.gov’.” 108NC explained, “I didn’t recognize the CorpVM email. I didn’t have a voicemail [sic] on my phone.” 105NC described, “I did not recognize the sender or the company, my voicemail is not connected to the web as far as I know.” 106NC elaborated, “I am not familiar with the company, it was from an unknown caller, I hovered over the hyperlink and the url looked odd and the voicemail light was not lit up on my phone. All very suspect!” 111NC explained, “I didn’t get a phone call prior to receiving the email. Made the email about voicemail very suspect.” 113NC reasoned, “I’ve never seen a message telling me I had a new voice mail,” and 101NC said, “Did not look like the normal email that I receive when I get a voice mail message.”

2) *Second exercise: unpaid invoice (August)*

As in the first exercise, responses from non-clickers differed greatly from those of clickers. For most respondents whose work responsibilities included handling financial matters—such as paying vendors or dealing with contracts and grants—the email was particularly believable and concerning given their individual user context. This was evident in clickers’ comments on their initial impressions of the email and their decisions to open the attachment. For example, 202C reported thinking, “That I had an unpaid invoice from when I assisted my division with credit card orders,” and 203C said, “that I missed a payment on one of my contracts.” 209C was concerned, “I had forgotten about an invoice for either a contract or purchase order.” 207C questioned, “What did I order that I haven’t paid for?” and 210C noted, “I pay invoices so I was wondering what invoice this was that did not get paid.”

Clickers also referred to recent real-world events that were congruous with the unpaid invoice email. 204C explained, “I had just talked to someone in accounts payable the previous day who told me they were new. I thought they were reaching out by email this time bc I saw the .gov” and 208C said, “We recently had a legit email from a vendor regarding unpaid invoices so I thought this may be another one. I work with private vendors so it looked legit.”

The fictitious sender’s @nist.gov email address was a particularly believable cue for clickers, with comments like “it came from a Fed” (203C) and “The email was from an internal email address.” (207C) Similarly, 210C reasoned, “From a NIST employee, figured she worked in AR and/or finance.” (AR stands for Accounts Receivable)

In general, non-clickers found the premise of this email particularly suspicious given their user context: either their work responsibilities did not include finances, or if they did, they found enough suspicious cues to prevent them from opening the attachment. For non-clickers, the mismatches with user context and expectation were particularly salient in this exercise, triggering enough suspicion to prevent them from opening the attachment. 203NC said, “My first impression was to ask why I was getting an email regarding payment of an invoice because I don’t have a credit card or purchasing authority. It made me suspicious.” 204NC elaborated, “I immediately thought this seemed like a strange email to receive. I don’t usually get emails about invoices that need to be paid, and certainly not from a NIST email address. I also thought it was strange that the attachment was a zip file.” 206NC noted, “I definitely thought it was suspicious since I don’t get any invoice related stuff.” 207NC said, “I don’t know this person and if they are from NIST they should not be sending an invoice this way.” 212NC explained, “I don’t deal with invoices or anything having to do with accounts payable or accounts receivable and haven’t recently purchased anything.” 213NC said, “The email did not match something that a federal employee would write. It seemed more like an email from a vendor.”

Importantly, although some non-clickers had finance-related responsibilities, they still found enough suspicious cues that they chose not to open the attachment. For instance, 202NC said, “The format of the email did not match the format of other invoices sent by accounts payable. The invoice number did not match the series used in either of the projects that I manage. The email was also addressed to my last name (not the usual format). The attachment was a zip file. Invoices are never in zip files.”

Like clickers, non-clickers also noticed and referred to the .gov email address cue. However, they performed additional fact-checking that clickers did not: they searched for the fictitious Jill Preston in the employee directory. 204NC said, “I searched for the sender in the directory and could not find them.” 205NC explained, “I looked up Jill Preston in the user directory and the person didn’t exist.” 206NC said, “I looked up the name in [the email client].” 207NC noted, “Looked the person up in the directory and did not find them.” 213NC noted, “I searched for the person in the phone directory and did not find anyone.”

In a quote that perfectly illustrates a non-clicker’s progression from general suspicion, to noticing a specific cue, to engaging in fact checking, and finally to reporting, 201NC explained, “...upon re-reading the email I became very suspicious. The email references a .doc attachment, but the attachment was a .zip file. After noticing that, I checked the NIST directory and saw that there was not a Jill Preston (Fed) at NIST. I immediately forwarded to my ITSO.”

3) *Third exercise: order confirmation (December)*

Although this exercise had the lowest click rate of the three, the importance of a user’s individual work context was again clear. When the premise of the email seemed believable given

user context, respondents clicked on the link. This match between an individual's reality and the email was clear for clickers. 302C explained, "We have some items on back order so I thought that this may be one of those so I clicked on it to see what the item was." 303C said, "I have several orders open and sometimes get email shipping confirmations," and "I did get a legitimate one that day as well."

In contrast, the premise of the email did not match reality for non-clickers' work contexts. For instance, 319NC explained, "This is weird because I don't place orders or pay for them." 305NC said, "This looks like a personal order. I never use my work email for ordering personal items." 310NC noted, "Suspicious since I didn't remember ordering anything." 330NC said, "Became cautionary after realized that this did not fit anything I had recently purchased."

As in the preceding two exercises, non-clickers focused more on suspicion-raising cues and reasoning through the cues and email content when making their decisions not to click the link. 301NC explained, "After recognizing an amazon imitation, I looked at the email address 'discountcomputers.com'. Never heard of it, or bought anything from there." 310NC also referred to the email address: "Discount was misspelled [sic] on the sender's email address." Similarly, 304NC said, "I noticed that 'discount' was spelled incorrectly in the email address."

Non-clickers also referred to the lack of specificity in the email's content. For example, 320NC explained, "As an unknown subject from an unknown source, with rather vague & generic content, it was obviously not legit." 309NC said, "There was no information identifying the company that was sending the email and there was no product information." 305NC noted, "This does not look like it comes from our Ordering systems. It had no descriptive organization or company."

As this exercise took place in December, the overall holiday shopping context was important, something that came up in comments from both clickers and non-clickers. For example, 309NC said, "I have been ordering some Christmas stuff but I have always used my home email so I was a little confused as to why this was coming to my work email."

B. Themes across exercises

1) User context

Across all three phishing exercises, a user's individual work context was key in understanding an individual's click decision. The importance of alignment/misalignment with a user's work context, expectations, and external events was clear in each exercise. Whereas clickers found the premise of the phishing emails to be particularly believable given their user context, non-clickers did not. Clickers tended to focus on compelling cues, such as the basic premise of the emails, or the .gov email address. However, non-clickers focused on suspicious cues like misspellings and unexpected attachment types. Because their user context led them to question the premise of an email, non-clickers reported performing additional fact-checking, such as searching for the sender in the employee directory.

For clickers, the emails were plausible enough given their work context that deeper thought and analysis were not triggered, indicating more surface level thinking. Even if something did seem out of the ordinary to them, it was insufficient to trigger alarm bells, or any suspicion raised was

counteracted by other believable cues, such as the presence of a .gov email address. For instance, 201C said, "Seemed a little out of the ordinary but legit email address." 104C said, "thought it was odd, but didn't connect the dots between phishing scheme and voicemail." In contrast, for non-clickers, the emails were suspicious enough that deeper thought was triggered: they questioned multiple cues, reasoned carefully about the emails, and even took additional steps to check facts when possible. Non-clickers also reported re-reading emails, for instance, 201NC said, "Checked address and saw (FED) so automatically assumed it was legit. After reading email, it was a little off, so I reread a few times."

2) Strategies and fact-checking

Across exercises, non-clickers described specific strategies they used and cues they looked for. For example, 326NC said, "I check multiple items, the from, the address when you mouse over, the content in general, look for misspelled words, strangely worded emails." 101NC described, "I have certain rules that I follow, I don't click on any links when I don't know who the email is from, and I don't click on any links when the email is from someone that I know but there is no message from that person." 325NC elaborated, "It reminds me of what to look for like a peculiar email address, odd salutation, bad grammar, and a sense of urgency in the message to make you do something that you normally wouldn't do. Because of this training, I am highly suspicious and my 'Spidey Sense' tingles whenever I see one of these emails."

In addition to following strategies and behavioral rules, non-clickers also engaged in fact-checking tailored to the premise of the phishing exercise. In the first exercise, they checked the voicemail light on their phone. In the second exercise, they searched for the fictitious Jill Preston in the employee directory. In the third exercise, they thought back through whether they had recently ordered anything. There were also several unique fact-checking strategies reported. For the second exercise, 103NC explained, "I checked the serial number of my work [mobile phone] to see if it was the same number as in the email." (Presumably this person was referring to the "receiving machine's ID" listed in the exercise email). Additionally, 304NC reported a clever fact-checking strategy regarding the sales tax listed in the third exercise email: "The MD sales tax is 6%. I calculated the tax based on the cost (\$59.97), and the sales tax listed in the email is greater than 6%."

While clickers did not report these types of fact-checking strategies in relation to the phishing emails, several engaged in fact-checking regarding the survey itself: for the second phishing exercise, three participants called the NIST researchers to verify that the survey was not a phishing attempt. When speaking with us, they said they had just been caught by the training and wanted to check this was not another phish.

3) Concern over consequences

Concern over potential consequences differed greatly between clickers and non-clickers. Across the three exercises, clickers were often concerned over consequences or repercussions arising from *not* clicking, such as failing to act or failing to be responsive. Referring to the voicemail phishing exercise, 105C explained, "I am always interested in ensuring that I get any messages and act on them. It could have been my supervisor or other person requiring an action on my part."

Clickers were particularly concerned over the consequences of not addressing an unpaid invoice: if they did not click to open the attachment, they could not pay the invoice nor could they forward it to the correct person for payment if it was not their direct responsibility. The unpaid invoice exercise also had the highest click rate, nearly double that of the other two exercises. The fact that an actual unpaid vendor invoice had recently been an issue in that OU no doubt contributed to clickers' increased concern over the email. 201C referred to "repercussions for not remitting payment," 203C responded, "potential to miss contract payments," 207C lamented, "I would have an unpaid invoice," and 208C worried, "I may miss a legitimate complaint/issue." 208C explained, "If true, I would be the person who would have to address the issue." 202C noted, "In the past 10 months I had 2 cancelled orders and wondered if it had gone through."

Across all three exercises, non-clickers were more concerned with consequences that could arise from clicking, such as downloading malware and viruses. For example, 103NC said, "I was concerned something might be downloaded onto my computer or I could get a virus." 106NC explained, "If this was a phishing email, I could be the person that allows someone access and they could potentially infect or steal information." 204NC said, "I considered that it might contain a virus or other malware." 210NC said, "I considered a hacker planting a listening bug onto the NIST systems." 211NC explained, "It did not look like a legit email so clicking on the link could have lead [sic] to virus, spyware, malware, etc." Similarly, 306NC said, "I thought if I clicked I could get a virus or have malware put on my computer."

4) *Confidence in institutional computer security measures*

Clickers seemed quite confident in the institution's security measures. For example, 101C stated, "I thought NIST security system can filter phishing emails." 105C said, "We are within a firewall at NIST?" 303C noted, "I do not get spam or junk emails (very often if at all), which tells me NIST is very pro-active in stopping them before we get them."

Non-clickers seemed to have a more tempered view, frequently referring to the idea that some phishing emails will always get through the filters. For example, 201NC explained, "Because of the widely varied nature of the work here, I am not sure it is possible to block out all phishing attempts." 202NC said simply, "It's impossible to catch them all." 213NC explained, "My feeling is that some phishing emails will get through no matter how good the security measures are." 320NC, "I think NIST puts serious effort into computer and internet security but things are always going to get through on occasion." 312NC noted, "There is no perfect security. The best NIST can hope for is to mitigate the number of attacks." 324NC said, "the attacks are constant, some will get through." Interestingly, 203NC referred to progress for both the institution and the hackers, "NIST has made a lot of progress in stopping phishing emails, and has trained staff well. However, I know that hacking continues to advance, too."

V. DISCUSSION

A. *Benefits of operational data*

The benefits of long-term, in situ operational phishing data are many. Specifically, the benefit of ecological validity is extremely valuable. By conducting an assessment using a

sample of real-world staff working in their normal work environments, we were able to obtain a high degree of ecological validity lacking in many assessments. Importantly, our sample was varied in terms of age, education, supervisory experience, and years spent working at the institution. Additionally, we surveyed exercise participants almost immediately after they clicked, in contrast to [4], where many months elapsed between the time participants received the initial phishing and were interviewed about it. We also addressed the "adversary modeling challenge" identified by Garfinkel and Lipford in [9], by using phishing emails modeled on real-world threats current with each exercise. As the nature of phishing continues to change, it is imperative to model attacks after real-world threats.

We believe that our long-term, in situ assessment with a varied sample should be meaningful and compelling for usable security researchers and security practitioners alike. Our workplace setting allowed us to both confirm findings from laboratory settings and unearth new findings that were only possible with a real-world setting.

B. *User context: an explanatory variable in phishing susceptibility*

A user's context—in this case, their individual work context—is the lens through which they interpret both the general premise of an email, as well as specific cues within the email. Through three phishing awareness training exercises and corresponding surveys, conducted in situ over a 10-month period, we found clickers and non-clickers interpreted the premise of phishing emails very differently depending on their individual work context. When a user's work context aligned with the premise of the email, they tended to find the premise believable and attended to compelling cues. In contrast, when a user's work context was misaligned with an email's premise, they tended to find the premise suspicious and they instead focused on specific suspicious cues.

Once we identified user context as an explanatory factor, we then saw that user context alignment coincided with depth of processing and differences in concern over consequences between clickers and non-clickers, which are discussed next. From respondent feedback, the initial read of the email by clickers seemed to produce a reaction that aligns with findings of Wang et al., where "Visceral triggers reduce the recipients' depth of information processing and induce recipients to make decision errors" [27] and Kumaraguru et al. in [15]. On the other hand, non-clickers reported a reaction of suspicion to the premise misalignment with their user context, helping them attend to phishing deception cues. We also discuss the operational setting and user context.

1) *Depth of processing*

Clickers and non-clickers seemed to process the phishing awareness training emails quite differently. When the general premise of a phishing email was believable for clickers, they did not then engage in deeper processing; they did not look for deception indicators, nor did they question the email's legitimacy. However, we saw evidence of extremely efficacious cue detection and utilization strategies on the part of non-clickers, such as checking for misspelled words, grammatical errors, mismatches, and so on. Our data also show that non-clickers report engaging in additional fact-checking behaviors, such as searching for the sender in the employee directory.

From a cognitive science perspective, this makes perfect sense: why waste effort and time on double-checking something that seems perfectly legitimate? In contrast, when the general premise of a phishing email was suspicious for non-clickers or they attended to specific suspicious cues, they engaged in deeper processing to confirm that the email was indeed a phish. For example, non-clickers reported re-reading the emails and engaged in additional fact-checking to follow-up on suspicious cues, such as checking their phone lights and searching for Jill Preston in the institution's employee directory. This difference in surface versus deep processing ties back to System 1 and System 2 thinking; with System 1 being fast, intuitive, and emotional, relying heavily on habits and heuristics, and System 2 being slower, more deliberative, and more logical [12]. This is in no way to say that clickers and non-clickers differ in depth of processing in general, only that for these particular phishing awareness exercises set in the real world, their survey responses showed evidence of differential processing.

2) *Consequence considerations*

Overall, both clickers and non-clickers provided feedback that conveyed an attitude of conscientiousness in the workplace. Despite that general conscientiousness, concern over consequences typically aligned with the click decision. Overall, clickers tended to be more concerned with potential consequences that could arise from *not* clicking: failing to act, seeming unresponsive to an email, or not addressing a legitimate issue. This was especially true for clickers in the unpaid invoice phishing premise, and seemed exacerbated by recent events where an unpaid invoice had been an issue for the OU participating in the awareness training exercises. In contrast, non-clickers were concerned with potential consequences that could arise from clicking: malware, viruses, and so on.

3) *The operational setting and user context*

It is important to realize that despite being part of the same approximately 70-person OU within the institute, staff in our assessment had very different *individual* work contexts. Responsibilities within the OU ranged from financial matters, such as processing orders, invoices, grants, contracts, and payments; to organizational safety matters relating to radiation, health and environmental issues; to program development and compliance. Even within the subset of staff members who shared somewhat similar responsibilities, individual work contexts varied. For example, within the subset of staff who had financial responsibilities, individual contexts varied depending on whether someone was recently working on a particular invoice and with whom they were working.

In addition to an individual's work context, another work context is the shared work context. For example, consider the collective awareness among our assessment participants regarding a recent unpaid vendor invoice. This shared awareness of a recent workplace issue affected participants' concern over the consequences of an unpaid invoice and may have contributed to elevated click rates for that phishing exercise. Though not uniquely a work context, an additional example of a shared context was the holiday context in the third exercise—that exercise leveraged the fact that more people place online orders near the holidays.

Such variability in individual work contexts, that at times partially overlaps with others' work contexts, is likely the case

at other institutions as well. It would be rare to find an organization where every member of the staff has exactly the same, or even completely different, tasks in at least medium-sized or larger organizations. Different people will be more or less susceptible to a particular phishing email's premise based in part on their individual work environments, tasks, and responsibilities.

C. *Revisiting prior quantitative data*

Based on the primary contribution from our 10-month operational assessment—showing the importance of individual user context in explaining phishing email click decisions—we are now able to better interpret the previously puzzling variability in click rates observed across the initial 3.5 years of phishing awareness training exercises at a U.S. government institution. Although we do not have user feedback from earlier exercises, the lens of the user context and the premise of each prior exercise together give new insight in understanding click results. Fig 1. shows click-rate data. Several data points deserve special mention. The exercise labeled 'From Rich' had a high click rate (49.3 %) and used a message that was tailored to mimic an actual spear phishing attack on the organization just prior to the exercise. The premise in this exercise aligned extremely well with the communication practices of the organization and the culture of conscientiousness. The phish labeled 'Web Logs' purported to be an automated message citing a violation of the institute's internet policy; it also had a high click rate (43.8 %). The premise this time aligned with the fact that the organization had such a policy and staff were aware of it. The scenario with the lowest click rate, 1.6 %, is labeled 'Ebola Outbreak.' Its low click rate is likely due to the circumstance that the institute's spam filtering placed the training phish in people's spam folder; despite this, one person saw it there and clicked its attachment. Knowing that people in this OU were responsible for organizational health issues at the institute and noting that the timing coincided with the Ebola crisis of 2014, we better understand the premise alignment. Click rates from the other exercises during 2012 to 2015 are between these extremes; each had plausible premise alignment with some staff members' job responsibilities within the OU or an external event that was relevant to the individual. This provides a plausible explanation why the institute's ITSND observed such intriguing variability in click rates across the initial 3.5 years of their trial deployment.

D. *Findings with actionable operational implications*

With data from an operational setting, we have the rare opportunity to make novel observations that have operational implications as well as confirm findings from other studies that hold in a new operational environment. We report on three findings that are immediately actionable by security practitioners.

1) *Setting realistic expectations of institutional security*

We found that participants in our assessment, especially clickers, expressed a great deal of confidence in the institution's technological security measures, perhaps too much confidence. While it is good that staff are aware there are indeed institutional security measures in place, having too much faith in such mechanisms can be dangerous if it leads staff to a false sense of total security. If people are overconfident in, or overly reliant on, institutional security measures, this may have an undesirable

effect on their clicking behaviors. They may not be concerned about clicking if they believe they are already fully protected. Organizations should consider explicitly educating their employees that no technological solution is completely infallible, especially reactive ones like those still commonly in use. Our operationally-situated findings of staff confidence in the institution's security measures provide confirmation to those described in [6] and [7], both of which were performed in other settings.

2) *The changing nature of phishing attacks*

Our exploration of click decisions revealed that some members of the sample did not know that the nature of phishing has evolved beyond the traditional requests for sensitive information such as usernames and passwords. Others, such as [7], have reported that users who know the definition of phishing are statistically less likely to fall for phishing than those who do not know the definition. This highlights the operational opportunity to ensure staff remain aware of how phishing attacks continue to evolve.

3) *Gamification*

Our ITSO partner observed emergent competition or *gamification* of the phishing awareness training exercises over the years, where people would try to beat their colleagues and be the first person to spot the training phishing email. In an effort to have each member in the OU make their training click decisions without co-worker aid, the ITSO cautioned people not to warn others of the emails. However, friendly competition among colleagues may be an unintended benefit to conducting long-term phishing awareness training exercises. Since phishing messages missed by technological solutions are often found by individuals reporting suspected phish or after negative consequences arise, an increase in the number of those reporting phishing messages, especially early reporting, should provide meaningful benefits to individuals and the information technology security incident-response teams tasked with detecting, containing, eradicating, and recovering from infected machines and networks [5]. Organizations have the opportunity to use that friendly competition to encourage reporting suspected phishing messages.

E. Implications for predicting phishing susceptibility

Our data establish user context as an explanatory variable in phishing susceptibility. Given the long-term, in situ nature of our data across a wide range of ages, education levels, and work experience, we believe our data are compelling and comprehensive. Of course, all new findings should be validated by additional studies, as should this one. However, studies that attempt to examine user context must be set in participants' real-world settings; laboratory settings cannot provide a sufficient level of ecological validity for this particular phenomenon. Additionally, models that attempt to predict phishing susceptibility such as those documented in [23] and [26] should further explore the user work context identified in this paper. This is in addition to other phishing susceptibility factors such as personality and individual risk tolerance, which others have studied, for example [3], [7], [10], [14].

VI. CONCLUSIONS

The problem of phishing is not yet solved and the impact on individuals and organizations continues to grow as attacks

become increasingly sophisticated. We have uncovered another piece of the phishing puzzle through an examination of long-term, operationally-situated data captured during embedded phishing awareness training exercises conducted at a U.S. government institution. In this paper, we focus on three exercises conducted in 2016, each having participant feedback. For all 15 exercises in the 4.5-year span (2012 through 2016), the phishing training emails modeled real-world phishing campaigns, and participating staff were in their normal work environments with their regular work loads, providing ecological validity.

This data has revealed the novel finding that the alignment of user context and the phishing message premise is a significant explanatory factor in phishing susceptibility, impacting both an individual's depth of processing and concern over consequences. Additionally, our data provide an operationally-relevant perspective from users regarding other factors affecting their click decisions. We believe that the rare opportunity to collect data surrounding phishing click decisions in the workplace coupled with ecologically valid conditions over time provided the crucial elements for these findings to surface. Because of the scarcity of operationally-situated data, we also report on findings we believe to be actionable now in operational settings. Our long-term operational data offer significant contributions to the existing corpus of phishing literature with implications for both future research and operational security.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the institution's ITSND and our ITSO partner for their support throughout this project, and thank anonymous reviewers for their comments.

DISCLAIMER

Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by the National Institute of Standards and Technology nor does it imply that the products mentioned are necessarily the best available for the purpose.

REFERENCES

- [1] G. Aaron and R. Rasmussen, "Global phishing survey 2016: Trends and domain name use," Anti-Phishing Working Group, June 2017. https://docs.apwg.org/reports/APWG_Global_Phishing_Report_2015-2016.pdf (Accessed Aug 2017).
- [2] A. Almonnani, B.B. Gupta, Samer Atawneh, A. Meulenberg, and E. Almomani, "A survey of phishing email filtering techniques," IEEE Communications Surveys & Tutorials, vol. 15, no. 4, Fourth Quarter 2013.
- [3] M. Blythe, H. Petrie, and J.A. Clark, "F for fake: Four studies on how we fall for phish," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, May 2011, pp. 3469-3478.
- [4] D. Caputo, S.L. Pfleeger, J. Freeman, and M. Johnson, "Going spear phishing: Exploring embedded training and awareness," in IEEE Security & Privacy, 12(1), January 2014, pp. 28-38.
- [5] P. Cichonski, T. Millar, T. Grance, and K. Scarfone, "Computer security incident handling guide," NIST Special Publication 800-61, rev. 2, 2012, <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf> (Accessed Oct 2017).

- [6] D. Conway, R. Taib, M. Harris, K. Yu, S. Berkovshky, and F. Chen, "A qualitative investigation of bank employee experiences of information security and phishing," in Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017), USENIX Association, 2017, pp. 115-129.
- [7] J.S. Downs, M. Holbrook, and L.F. Cranor, "Decision strategies and susceptibility to phishing," in Proceedings of the Second Symposium on Usable Privacy and Security (SOUPS '06), ACM, pp. 79-90.
- [8] A.J. Ferguson, "Fostering e-mail security awareness: The West Point carronade," *EDUCASE Quarterly*, 2005, 1, pp. 54-57.
- [9] S. Garfinkel and H.R. Lipford, Usable security: History, themes, and challenges. Synthesis Lectures on Information Security, Privacy, and Trust, E. Bertino and R. Sandhu, Eds., 5(2), Morgan and Claypool Publishers, 2014, pp. 4-6, 55-65.
- [10] T. Halevi, N. Memon, O. Nov, "Spear-phishing in the wild: A real-world study of personality, phishing self-efficacy and vulnerability to spear-phishing attacks", January 2, 2015, <https://ssrn.com/abstract=2544742> or <http://dx.doi.org/10.2139/ssrn.2544742> (Accessed Oct 2017).
- [11] J. Hong, "The state of phishing attacks," *Communications of the ACM*, 55:1, January 2012, pp. 74-81.
- [12] D. Kahneman, *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.
- [13] K.L.K. Koskey, T.A. Sondergeld, V.C. Stewart, and K.J. Pugh, "An applications of the mixed methods instrument development and construct validation process: Transformative experience questionnaire," *Journal of Mixed Methods Research*, 2016, 1558689816633310.
- [14] P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L.F. Cranor, and J. Hong, "Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer," in Proceedings of the anti-Phishing Working Groups 2nd Annual eCrime Researchers Summit, (eCrime '07), ACM, October 2007, pp. 70-81.
- [15] P. Kumaraguru, S. Sheng, A. Acquisti, L.F. Cranor, and J. Hong, "Lessons from a real world evaluation of anti-phishing training," in e-Crime Researchers Summit, IEEE, October 2008, pp. 1-12.
- [16] P. Kumaraguru, J. Cranshaw, A. Acquisti, L.F. Cranor, J. Hong, M. Blair, and T. Pham, "School of phish: A real-world evaluation of anti-phishing training," in Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS '09), ACM, July 2009, article 3.
- [17] New York State Office of Cyber Security & Critical Infrastructure Coordination, "Gone phishing... a briefing on the anti-phishing exercise initiative for new york state government: Aggregate exercise results for public release," 2005. (Not available online. Additionally, the NYS Office of Information Technology Services no longer has the report—contacted 8/28/2017).
- [18] J. Nicholson, L. Coventry, and P. Briggs. "Can we fight social engineering attacks by social means? Assessing social salience as a means to improve phish detection," in Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017), USENIX Association, 2017, pp. 285-298.
- [19] J. Saldaña, *The Coding Manual for Qualitative Researchers*, 2nd ed., Sage Publications, 2013.
- [20] SonicWall, "2017 SonicWall annual threat report," 2017. <https://www.sonicwall.com/docs/2017-sonicwall-annual-threat-report-white-paper-24934.pdf> (Accessed Aug 2017).
- [21] Stanford University, "University IT launches phishing awareness service," 2016. <https://uit.stanford.edu/newsletter/university-it-launches-phishing-awareness-service> (Accessed Aug 2017).
- [22] Symantec, "Internet security threat report," vol. 22, April 2017. <https://www.symantec.com/security-center/threat-report> (Accessed Aug, 2017).
- [23] F.P. Tamborello, K.K. Greene. "Exploratory Lens Model of Decision-Making in a Potential Phishing Attack Scenario." National Institute of Standards and Technology Interagency Report, NISTIR 8194, October 2017, DOI: <https://doi.org/10.6028/NIST.IR.8194>
- [24] C. Teddlie and A. Tashakkori, *Foundations of Mixed Methods Research: Integrating Quantitative and Qualitative Approaches In The Social and Behavioral Sciences*. Thousand Oaks, CA: Sage, 2009.
- [25] Trend Micro, "Spear phishing 101: What is spear phishing?," September 2015. <https://www.trendmicro.com/vinfo/us/security/news/cyber-attacks/spear-phishing-101-what-is-spear-phishing> (Accessed Aug 2017).
- [26] A. Vishwanath, T. Herath, R. Chen, J. Wang, and H.R. Rao. "Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model." *Decision Support Systems*, vol. 51 no. 3, March 2011, pp. 576-586.
- [27] J. Wang, T. Herath, R. Chen, A. Vishwanath, and H.R. Rao, "Phishing susceptibility: An investigation into the processing of a targeted spear phishing email," in *IEEE Transactions on Professional Communication*, vol. 55, no. 4, December 2012, pp. 345-362.

APPENDIX

This appendix contains a detailed description of the survey instrument template, including customizations, questions, response options, and show/hide question logic for follow-up questions. Survey questions are shown in Table I below. Gray highlighting is used to indicate page breaks in the survey.

A. Landing page

Each survey started with a landing page that briefly described the nature of the survey, the estimated time it would take a respondent to complete the survey, and the potential risks and benefits associated with taking the survey. The landing page also contained the researchers' contact information.

B. Customizations

As described in the Method section, the survey template was created such that it could be easily customized based on the nature of the phishing exercise. After the landing page, the first page of the survey was customized by adding a screenshot of the phishing email at the top. Open-ended questions about initial impressions (Q1) and click decisions (Q2) were together on the first page. Q2 was customized depending on whether the phishing email contained a link or an attachment, and depending on whether a respondent was a clicker or a non-clicker. Q2 customizations were as follows:

[Clickers, link] What made you decide to click the link?

[Non-clickers, link] What made you decide not to click the link?

[Clickers, attachment] What made you decide to open the attachment?

[Non-clickers, attachment] What made you decide not to open the attachment?

TABLE I. SURVEY QUESTIONS.

Q#	Primary question			Follow-up question			
	Question	Response type	Response options	Show/hide logic	Question	Response type	Response options
1	Think back to when you first saw the email above in your inbox. What were your initial impressions?	Open-ended		n/a			
2	What made you decide not to click the link? Note that Q2 was customized for link/attachment and clicker/non-clicker (as described above).	Open-ended		n/a			
3	What type of device were you using when you opened the email?	Radio buttons	Desktop computer Laptop Tablet Smartphone Don't remember Other – Write In: _____	n/a			
4	Were you in a hurry when you opened the email?	Radio buttons	No Yes Don't remember	n/a			
5	How thoroughly, if at all, did you read the email?	Radio buttons	Didn't read it Read some of it Read it quickly Read it thoroughly Don't remember	n/a			
6	What were you doing around the time you opened the email?	Radio buttons	Focused solely on checking email Checking email while primarily focused on another task Don't remember Other – Write In: _____	n/a			
7	Did anything in the email seem pertinent to you?	Radio buttons	No Yes Not sure if pertinent Don't remember				
				Shown only if "yes" or "not sure"	In what way?	Open-ended	
8	Did you consider any possible consequences for clicking on the email link?	Radio buttons	No Yes Not sure				
				Shown only if "yes"	What possible consequences did you consider?	Open-ended	
				Shown only if "yes"	How concerned were you about these consequences you listed?	Radio buttons	Not at all concerned Somewhat concerned Very concerned Don't remember if I was concerned at the time
9	Did you consider any possible consequences for <u>not</u> clicking on the email link?	Radio buttons	No Yes Not sure				
				Shown only if "yes"	What possible	Open-ended	

Primary question				Follow-up question			
Q#	Question	Response type	Response options	Show/hide logic	Question	Response type	Response options
					<i>consequences did you consider?</i>		
				Shown only if "yes"	<i>How concerned were you about these consequences you listed?</i>	Radio buttons	Not at all concerned Somewhat concerned Very concerned Don't remember if I was concerned at the time
10	<i>How curious, if at all, were you about the email?</i>	Radio buttons	Not at all curious Somewhat curious Very curious Don't remember				
				Shown for all responses but "don't remember"	<i>Why or why not?</i>	Open-ended	
11	<i>When thinking about the email, how suspicious, if at all, were you?</i>	Radio buttons	Not at all suspicious Somewhat suspicious Very suspicious Don't remember				
				Shown for all responses but "don't remember"	<i>Why or why not?</i>	Open-ended	
12	<i>How confident are you in NIST's computer security measures to prevent phishing emails from reaching you?</i>	Radio buttons	Not at all confident Somewhat confident Confident Very confident				
				Shown for all responses	<i>Explain.</i>	Open-ended	
13	<i>Overall, do you feel like the phishing awareness training emails have changed your email behavior?</i>	Radio buttons	No Yes Not sure				
				Shown only if "yes"	<i>How has the phishing training changed your behavior?</i>	Open-ended	
				Shown only if "no" or "not sure"	<i>Why?</i>		
14	<i>About how many emails do you receive in a typical work day?</i>	Radio buttons	0 1-5 6-10 11-20 21-30 31-40 41-50 51 or more Prefer not to answer				
15	<i>About how many emails do you send in a typical work day?</i>	Radio buttons	0 1-5 6-10 11-20 21-30 31-40 41-50 51 or more Prefer not to answer				

Primary question				Follow-up question			
Q#	Question	Response type	Response options	Show/hide logic	Question	Response type	Response options
16	<i>What is your gender?</i>	Radio buttons	Female Male Other identification Prefer not to answer				
17	<i>What year were you born?</i>	Drop-down	Options ranging from 1931 to 1998 Prefer not to answer				
18	<i>What is your highest completed level of education?</i>	Radio buttons	Highschool or GED Associate's Degree Bachelor's Degree Master's Degree Doctorate Degree MD Other – Write In: _____ Prefer not to answer				
19	<i>How long have you been working at NIST?</i>	Drop-down	Options ranging from Less than 1 year to 50 years More than 50 years Prefer not to answer				
20	<i>Any additional comments on this survey or the phishing exercises?</i>	Open-ended					