How Non-Experts Try to Detect Phishing Scam Emails

Norbert Nthala
Media and Information
Michigan State University
East Lansing, MI 48824
Email: nthalano@msu.edu

Rick Wash
Media and Information
Michigan State University
East Lansing, MI 48824
Email: wash@msu.edu

Abstract—**Email remains one of the most widely used methods of communication globally. However, successful phishing email attacks and subsequent costs remain unreasonably high despite technical advances in defenses that limit phishing scams. In this paper, we examine human detection of phishing. We found that non-experts go through four different sensemaking processes to determine if an email is a phishing message; they use different knowledge and skills to become suspicious differently in each process. Additionally, non-experts rely on their social connections as an investigative tool to determine if an email is a phishing scam. We discuss the impact of our findings on phishing training and technology.**

I. INTRODUCTION

Email is one of the most commonly used methods of communication, especially in large organizations and for e-commerce. Over 3.9 billion people have email accounts worldwide, and collectively they send and receive over 290 billion emails per day [1]. One feature of email that has contributed to its widespread adoption is that people can communicate with strangers; email is one of the major methods that we have to communicate with businesses and people we don't know and haven't interacted with before. However, because email is a global system where anyone can communicate with anyone, malicious actors conduct phishing scams by sending emails that pretend to be something that they are not, and trick people into taking actions that they otherwise wouldn't. [2]

Phishing messages, and phishing emails in particular, are an attack vector that has caused a large amount of damage in society. Phishing emails have been used to steal large amounts of money [3], install ransomware [4], or simply to steal email contents that are later made public [5]. 32% of all corporate breaches in 2018 were due to phishing [6]. Spear-phishing – a variant where emails are custom targeted to the recipients – is used by 65% of groups doing targeted attacks, and is more commonly used than zero-day vulnerabilities (only 23% of such groups) [7].

Our society has a number of defenses that help limit phishing scams, but do not completely prevent them. These defenses come in a number of different forms. Technological defenses try to detect known features of phishing emails (like known malicious attachments or known malicious links) and block emails. Organizational defenses allow society to try to take down phishing websites or prevent phishing attacks from receiving payments. Socio-technical defenses combine the work of people and computers, for example by allowing end users to report phishing emails, which are then investigated and removed from other users' inboxes or blacklisted. And finally, there are human defenses, where the recipient of the email recognizes the email as dangerous and chooses not to act on it.

In this paper, we focus on this last category: human detection of phishing. We examine the process that typical email users follow when looking at an email to determine if the email is a phishing message. We sent a (fake) phishing message to 31 people, and then interviewed them a few days later to get a retrospective account of the processes they followed when they received that message. We describe, in detail, a four-stage process that users typically follow for email messages, and then describe multiple ways this process changes as users become suspicious of the email.

II. RELATED WORK

Our work builds on the work of Wash [2], who looked at how IT experts identify phishing emails. He interviewed 21 IT experts about instances when they successfully identified emails as phishing in their inboxes. Wash describes that when an email is received, experts treat it like any other email — the content in the email is taken at face value and the person tries to make sense of the email and figure out what it is asking them to do. As they do this, they notice discrepancies — things that "feel off" about the email. Eventually, something triggers the person to think that this email is not legitimate — that it might be a phishing email that is not what it says it is. At this point, they become suspicious and begin explicitly looking for things that can help them determine if the email is legitimate or not. These new pieces of information often allow them to conclusively identify the emails as phishing.

III. METHODS

Our goal with this study was to examine everyday email users – non-experts who send and receive email regularly, but don't necessarily have technical training or regular exposure to cybersecurity attacks. We recruited 31 non-IT staff employees at a large university to participate in our research study about email. Table I shows demographics of the participants.

Education	Age			Gender	
	18 - 34	35 - 44	45+	Male	Female
Some college, no 4-year degree	0	0	2	0	2
4-year college degree	4	2	1	2	5
Some postgraduate or professional schooling, no postgraduate degree	0	1	5	2	4
Postgraduate or professional degree, including master's, doctorate, medical or law degree	7	4	5	8	8

TABLE I
Demographics of our participants

We recruited participants by asking heads of organizational units to share our recruitment message with employees in their units. Our recruitment message invited employees to participate in a "research study that aimed to understand how people who are not experts in computer security make security-related decisions". We indicated that eligible subjects would be asked to attend an interview lasting about 120 minutes where a researcher would ask them various questions to understand how they had made security decisions in the past or how they would make decisions in some given situations. Those interested were asked to contact us by email. In addition to advertising through heads of orgnizational units, we also advertised our study to employees of the university that are registered on SONA[1] — a paid system for recruiting study participants.

When a potential participant was interested in participating, they filled out a brief survey that included a consent form, basic demographic information, and questions about their use of a couple of common technologies used by the organization. We then asked them to schedule a 2-hour interview via Zoom videoconference. All interviews were scheduled at least two days in the future.

Approximately 2-3 days before each interview, we composed and sent a phishing email pretending to be from the organization's IT group informing them about changes to technology support (see sample email in appendix A). We always informed them about a technology they indicated they used most in the initial survey. The email included a link to a phishing website hosted at an off-campus, non-official domain. The website first presented the user with a fake login screen for the university, and then redirected them to a Google form based survey after they entered credentials on the login page. We recorded whether they tried to log in to the webpage (and how many times), but did not record any passwords entered or validate any passwords. This process was reviewed and approved by the university's IT department.

At the beginning of each interview, we reminded the subject about the goals of the study and their rights as a research participant. We asked them to consent again by completing an online consent form hosted on Qualtrics. After the subject consented, we debriefed them. We showed the subject the phishing email we had sent and asked them to remember what happened. We explained to the participant that we sent the email so that we could use it in our conversation to understand the process they followed when dealing with potentially dangerous or suspicious emails. The focus on the

process ensured that focus was not on whether the participant had fallen for our phishing email or not. After debriefing, we asked the participant if they wanted to continue participating in the interview. All participants accepted and participated in the interview.

During the 2-hour interview, we followed the Critical Decision Method [8] that was also used by Wash [2] for his interviews of experts. In the first phase, we used a hand-written shared display (a whiteboard using the Notes App on iPad) to draw a timeline of events that the participant described. This helped us to coordinate our understanding of the sequence of events related to the email. We made sure the subject was in control and directed the process when drawing the timeline of events. In the second phase, we went back through the timeline, and at each step, we asked a series of deepening questions intended to better understand what the participant was thinking: what their goals were, what options were considered, what background knowledge was used, and how they made decisions about dealing with the email. Finally, if time permitted, in the third phase, we made another sweep through the timeline and asked "What if?" questions, asking the participant what they would have done differently if parts of the email or event had been different. We avoided judgemental language througout the interview. After each interview, the interviewer wrote a short story description of the incident re-organizing the incident in the order in which events happened.

Participants were compensated with a $40 USD Amazon gift card for their time. The approach of the study and the interview protocol were reviewed and approved by our university's IRB as exempt.

We analyzed the data using open coding [9]. We used the short stories and interview transcripts to identify components of decision-making that each participant mentioned, including cues they noticed, actions they took, information they sought, sources of the information, and goals they were trying to achieve. After identifying these, we organized them in chronological order based on the timeline diagrams and compared across participants to identify patterns in the decision process. From this, we identified a four-stage process that non-experts follow to deal with emails, and how they become suspicious and investigate emails.

IV. DETECTING PHISHING EMAILS

A. Normal Process for Dealing with an Email

As non-experts receive emails, we found that they went through a four-stage process to deal with each incoming email.

At each stage, the person engaged in a sensemaking exercise; they had a specific goal they were trying to achieve and were trying to build an understanding of the situation around the email in order to achieve that goal. Each stage involved just enough sensemaking work [10] to accomplish the goal. Between each stage was a set of actions or decisions that allowed the person to move from one stage to the next.

The first stage of the process involved *getting context of the email*. During this stage, the person tried to understand how the email related to other aspects of their life (e.g. project, tasks, events), without engaging the email in any depth. The goal, at this stage, was to determine whether the email was relevant to them and assess the importance of the email.

In the second stage, *reading the email*, the person engaged the email at greater depth to understand it within the context identified in stage one. The goal of the second stage is to understand why they received the email, and what was being requested of them in the email.

Once the person had understood the email and figured out what to do, they proceeded to the third stage, *enacting the request*. The person started taking positive action(s) to fulfill the request in the email. As the person enacted the request, they also engaged a sense-making process to understand the request (as explained in the email) and the action they were taking, and how all these fit within the context.

In the last stage, *getting closure*, the person sought to move on from the email (tick a box that it was done). Sometimes, this involved re-reading the email and relating it to the completed task/action and the context at large. The core part of this stage was the feeling of a sense of closure, which was complemented by actions like deleting an email or marking an email as done.

We illustrate this process with a brief story from Alice[2] as she received a new email.

Case 1. Alice: It was on a Wednesday afternoon and Alice was at her desk going through emails in her work account. She used to sort her inbox alphabetically in ascending order and wanted to get through as many emails as possible.

She saw an email from the IT department of her organization that said "ACTION REQUIRED" in the subject line. She had received other numerous emails form the department before; they were sent from different email addresses/sender names which were all in her inbox at this time. She noticed that this email was on its own, and that she had never received an email from this address before. This made her question the legitimacy of the email.

However, she recalled that her department once asked the IT department to setup a unique email address to be used for one of their activities. This made her think that different departments across her organization had similar unique email addresses reserved for particular purposes (in this case, the IT department had this address for other purposes).

She decided to open and read the email to make sure she acted on what the university wanted her to act on. She clicked

on the email and started reading it in a preview window. She saw a link to a survey in the email which made her question the validity of the email further since she did not like to click on links in emails. In the past, one of her colleagues had clicked on a link in an email which infected her computer and those of others in the workplace to a point where the IT department had to replace some of the computers in the department. She feared the same would happen and was cautious about clicking on links in emails until she was sure the email was legitimate.

She noticed her department's name mentioned a couple of times in the email body, but also contact details for the IT department in the email signature. These made her feel that the email was more legitimate than she was suspicious of it.

She re-read the email to confirm validity of the email. She noticed the email had a line she identified as having used for a number of years when sending emails to others in her line of work. She questioned herself whether she had to send the email to her contact in IT for verification or not, or if she had to delete it. In the recent past, she had had an altercation with her boss about not acting on an email and she was afraid that if she did not act on this one on time, she would have problems with her boss again.

She decided to just click on the link and make sure she got everything done. The link opened a webpage with a survey. She felt the survey looked like something her organization would put together and send out. This made her feel confident that she had not clicked on a bad link. She completed the survey and marked the email as completed. She then moved on to do other things.

In this case, Alice went through all four stages for dealing with this email. She first tried to get the context of the email by understanding who it was from (the IT department). Once she was satisfied with this context, she clicked on it to read the email for more details. Once she understood what the email was asking, she acted on it; she clicked the link and filled out the survey. Finally, she got closure by marking it as completed and moving on to other things.

B. Instrumental Actions

During each of the four stages, participants would regularly take actions intended to further their goal of understanding the email. As these actions were part of the sensemaking task [11], participants founds them to be easy to recall. Sensemaking actions include clicking on the email to read it, or hovering over an abbreviated name to see the full name.

However, some of the actions they took during the process were instrumental rather than sensemaking; they were a necessary step along the way but weren't intended to provide new information. Instrumental actions were not integrated into the sensemaking process. People did not consider these actions to be significant, and indeed often were not able to remember doing them.

One particularly important instrumental action was logging in. In the transition from stage two to stage three, subjects would often click on the link in the email to get to the survey.

Along the way, they would be presented with a login page. Logging in was instrumental; since they weren't doing it for sensemaking purposes, it wasn't seen as important and wasn't integrated into their cognitive sensemaking activities. Often, subjects would not even remember logging in (even though we have log data that says that they did).

Out of 31 participants, we recorded URL clicks from 21 participants. 3 participants had reported the email to IT security, who clicked the links. 16 of the 21 participants tried to login on our phishing page after visiting our fake login page. 10 of the participants continued on to fill in the Google form that we presented after the fake login page. However, only 6 people remembered the login page during the interview.

For instance, after clicking the link and logging in, Charlie saw the Google form that appeared. He became suspicious of the email, and concluded that the email was not legitimate. He then thought back through the actions he could remember to see if he had done anything dangerous like opening an attachment or providing information. Even moments after the act, he did not remember logging in; it was instrumental to his main action (opening the link to see where it went) and thus was not memorable.

C. Becoming Suspicious and Investigating

As our participants went through the four stages to deal with the email, they might notice things that were weird or unusual. Sometimes, but not always, these unusual things would trigger the person to become *suspicious*. For us, "suspicious" is an explicit cognitive state where the person holds two competing ideas in their head at once: their existing understanding of the email, and an alternative understanding that the email is not what it says it is (i.e. that it is a phishing or scam email).

We found that it was possible for participants to become suspicious of the email at any stage of the process. Becoming suspicious was always part of the sensemaking work. Since they were trying to make sense of the situation, they were actively trying to integrate each new piece of information into their current frame of understanding of the email. As they did this, they would sometimes notice "discrepancies", or things that they could not integrate into their understanding. Upon noticing enough (or serious enough) discrepancies, they would sometimes become suspicious of the email. That is, they would identify a second, distinct possible explanation for the email that the email is not real, is not legitimate, and is not what it says it is. This process of becoming suspicious was similar to what Wash [2] found with experts. However, we identified that this "becoming suspicious" could happen for non-experts during any of the four sensemaking stages.

Once someone became suspicious, then they added a goal: to understand whether or not the email was real. To accomplish this, they would take explicit actions to investigate the email. These actions included re-reading the email to look for grammar or spelling mistakes, checking authenticity of contact details in the signature of the email, hovering over the link in the email, and checking the sender's email address. However, except the domain of the sender's email address, most of

these checked out which made it harder for participants to conclusively classify an email as real or not.

Some participants then turned to their colleagues to ask if they got the email; for instance, Bob thought his colleagues had received a similar email based on his understanding of the work context and he asked them directly:

> **Bob**: *"So if it was a legitimate policy, I would expect that other people within my group, within my office, would be receiving the same email... But because none of them had gotten it, it was wow, this is a really well put together phishing email."*

Some chose not to enact the request in the suspicious email until one of their colleagues talked about it; for instance, Clara chose wait to hear from colleagues to know if it was legitimate:

> **Clara**: *"So a lot of times I rely on kind of my connections too, to find out if... Like I said, if something is legit, I'm going to hear about it in some other way. Either a colleague is going to mention it to me, or I'm going to be in a meeting and someone's going to mention it."*

Some felt a duty of care towards others. They investigated who else might have received the suspicious email by asking their colleagues (via work groups on Microsoft Teams) or asking local IT if anyone had reported the email to them. Some would advise their colleagues not to act on the email if they received it until they finished investigating it.

Some reported to IT asking for their opinion/advice. This investigation was seen as very important; it would often override other previous goals until they figure it out.

V. BECOMING SUSPICIOUS IN THE FOUR STAGES

The normal process of dealing with an email involves four stages. We found that participants can become suspicious in three of those four stages, but that the way they become suspicious differs by stage.

A. Suspicion in the First Stage: Getting Context

In the first stage, participants expressed a goal of assessing the relevance and/or importance of the email. First, they needed to identify the context of their life that the email related to (work vs. personal, which work project?), so they could assess its relevance or importance in that context. During this first stage, participants felt time pressure, which led them to primarily use the stated identity of the sender and the stated subject of the email to identify that context. Once identified, they used this context to assess the relevance and importance of the email – is it worth reading? How urgently?

Participants indicated that for many emails, no context could be identified or the identified context was irrelevant to them. These emails were deleted, unread. This is also an effective phishing prevention strategy: this relevance judgement during the first stage likely eliminates many phishing emails without even determining whether they are phishing or not.

Most participants only looked at the sender's name to put an identity on the sender. If the name made sense to them

(fit within the context), they rarely checked the sender's email address ever after until after they were already suspicious. Our email bore the name "[ORG] IT Services", which (1) was a well known name for the department, and (2) followed the naming pattern used on most emails from the department.

Alice, however, became suspicious at this stage because she organized (or grouped) emails by sender. She noticed that despite receiving numerous emails from the IT department, this one did not group with any of those. She explained:

> *Alice: "The first thing I thought was there's other [ORG] IT emails in yet a separate, like it didn't fall under [ORG] IT services desk. It didn't fall under [ORG] IT security office. It did not fall under [ORG] IT announcement. It was just [ORG] IT services. It was all by itself. I had emails from [ORG] IT announcements and [ORG] IT security office and [ORG] IT services desk, but not [ORG] IT services. And that was, I believe that in that moment, that was one of the impetus in my brain for questioning the legitimacy of the email, because it just was all separate and on its own and not under other sender titles, I guess, for lack of a better description."*

Once the context of the email was determined relevant, our participants would move to the second stage: reading the email. However, they first had to decide whether to read the email immediately or delay until later. We suspect delays might impact how people become suspicious, but do not have definitive evidence yet.

B. Suspicion in the Second Stage: Reading Email

In the second stage, our participants would read the email in detail, trying to build an understanding of the email and its relationship to the identified context. They tried to understand how the email relates to their life. All of the phishing emails we sent indicated that there was an upcoming change in the employer-provided technology the recipients used. Our participants would think carefully about what that technology change would mean for their work and their work environment. For example, many participants were very concerned about changes to Zoom videoconferencing support since they were working at home during the COVID-19 pandemic.

As they did this sense making, they relied heavily on typicality recognition [12], [13]: they looked for aspects of this email that were similar to previous emails with a similar context. This provided an opportunity to notice "typicality violation" discrepancies [2], [14] — things about the email that were different than what was typically present in previous emails. Some participants questioned the presence of the link in the email, which they said was not present in previous emails from the IT department. Some participants had a rule of thumb "not to click on links" unless they knew the sender, or until they verified that it was a legitimate link.

Bob, for example, received an email informing him that old files would be archived or deleted from OneDrive. When he read the email, he noted that what the email said about files being archived or deleted from OneDrive due to a change in technology did not fit his understanding of how his organization made such changes:

> *Bob: "I read through what it said and what it was asking and what it was indicating the actions would be. And that was something that I hadn't heard would be happening and it sounded outside the norm of how archiving would work. Just because something's older doesn't mean it's unimportant."*

This logic led Bob to become suspicious and investigate the email further.

At this point, some participants would also take other sensemaking steps to understand the email. For example, the tone of our email made it sound like the email was sent to other people, which led some to contact others they knew who would have received it to ask them about it.

This sensemaking ended once the person felt like they understood what the email was asking for, and they knew what they needed to do in response to the email.

C. Suspicion in the Third Stage: Enacting the Request

In the third stage, our participants would take positive actions to enact the request in the email — filling out a form, in our case. As they enacted the request in the email, they engaged in a third sensemaking process to understand how they were supposed to do it. Our participants would try to understand what the survey was, how the results would be used, and how it related to the request in the email. They would also try to relate the request to previous occasions when they were asked to complete a survey by the organization.

Dave, for example, received an email about upcoming changes to Zoom and a request to provide relevant information to IT. Dave explained;

> *Dave: "I thought, because the form looked genuine because there was a [ORG] logo and the colors were [ORG colors], I looked at the questions and I started filling the form. And then, I started getting a little bit suspicious... Because they were asking about some personal information, like date of birth and address, which can be misused, right. And the forms that I have filled, the [ORG] forms which I filled in the past, they never asked for addresses or date of birth or probably cellphone number too. [...] So I think it's asking something which the previous form didn't ask. So I thought this is fishy because it's asking a lot of personal information."*

Seeing the survey, he became suspicious and investigated the email. He noticed the sender's email address was from a different domain than his organization. He then closed the tab that had the form open. Suspicion in this third stage still arises out of sensemaking, but sensemaking about accomplishing the task rather than about understanding the email.

This stage ended when the person felt that they had completed enacting the request.

D. Suspicion in the Fourth Stage: Getting Closure

There was one final sensemaking process that happened for each email. Almost all of our participants described a need for a sense of closure around the email — a feeling that they were done with the email. Until feeling the sense of closure, they would take actions to ensure they came back to finalize processing the email, such as leaving it in their inbox, marking it unread, or making a mental note to revisit the email.

Participants who had chosen not to do anything with the email after reading it mentioned still thinking about the email days later. They contemplated whether the email was legitimate and what consequences would follow if they did not do what the email asked. None of the participants who completed the form became suspicious during the closure phase.

However, if the person had become suspicious and had engaged investigative steps (e.g. reported the email to IT or asked colleagues), then they would wait until the investigation was completed before taking further actions. They did not feel like they had closure until they had confirmation of the (il)legitimacy of the email (e.g. until IT had responded).

VI. IMPACT ON SECURITY TRAINING AND TECHNOLOGY

We found that the process of becoming suspicious for non-experts is similar to what Wash [2] found for experts: in their sensemaking process, users notice enough discrepancies that trigger them to have a second explanation about the email — that it is not legitimate. In addition, our findings on the sensemaking process expound Wash's "sensemaking stage" into four separate sensemaking stages and describe how non-experts can become suspicious at any of these stages.

Second, current phishing training messages emphasize investigative steps that require an individual user to focus on features internal to their inbox to detect phishing: checking source or reply-to email address, verifying URLs, looking for typos, urgency, etc. [15]. We found that users also depend on their social connections (an existing and available resource in most work places) as an important investigative tool for detecting phishing, unlike experts who mostly rely on advanced technical investigations [2]. As phishing emails get complex and harder for most non-experts to detect, these social connections play a key role in combating phishing attacks. This finding is consistent with other work (e.g. [16], [17]) on the role of social navigation in security and privacy.

Third, we found that humans have multiple sensemaking processes (stage one through three) for detecting phishing. Training users on detection techniques in each stage separately can encourage a defense in depth approach to phishing detection.

Fourth, logging in was instrumetal and most participants did not remember doing it. This presents a challenge, especially when someone reports or deletes the email after enacting the request; they would not know the extent of the problem. This can be mitigated by adding features to email programs to track user actions (including those external to the inbox, e.g. logging in) and maintain an accessible log a user can access. This can help the user recall their actions and can be handy for incident response. This solution, however, can raise privacy issues. The field of privacy enhancing technologies (PETs) has matured over time and has developed various approacches and techniques of ensuring user privacy in different contexts [18]. We postulate that such techniques can be applied to solve this problem (e.g. anonymity of a user's specific actions to all but the user); but needs futher exploration to ascertain best ways to achieve this.

REFERENCES

[1] T. R. Group, "Email statistics report 2019-2023 executive summary," The Radicati Group, Tech. Rep., 2019.

[2] R. Wash, "How experts detect phishing scam emails," *Proceedings of the ACM: Human Computer Interaction*, vol. 4, no. CSCW, p. 160, October 2020.

[3] MacEwan University, "University Discovers Online Fraud," 2017. [Online]. Available: https://www.macewan.ca/wcm/MacEwanNews/PHISHING_ATTACK

[4] R. Smith, "How a U.S. Utility Got Hacked," *Wall Street Journal*, Dec 2016. [Online]. Available: https://www.wsj.com/articles/how-a-u-s-utility-got-hacked-1483120856

[5] E. Lipton, D. E. Sanger, and S. Shane, "The Perfect Weapon: How Russian Cyberpower Invaded the U.S." *The New York Times*, dec 2016. [Online]. Available: https://www.nytimes.com/2016/12/13/us/politics/russia-hack-election-dnc.html

[6] Verizon, "2019 Data Breach Investigations Report," Tech. Rep., 2019. [Online]. Available: https://enterprise.verizon.com/resources/reports/dbir/

[7] Symantec, "Internet Security Threat Report," Tech. Rep. February, 2019. [Online]. Available: https://www.symantec.com/security-center/threat-report

[8] B. Crandall, G. Klein, and R. Hoffman, *Working Minds: A Practitioner's Guide to Cognitive Task Analysis*. A Bradford Book, 2006.

[9] S. H. Khandkar, "Open coding," *University of Calgary*, vol. 23, p. 2009, 2009.

[10] K. E. Weick, *Sensemaking in Organizations*. Sage Publications, 1995.

[11] G. A. Klein, J. K. Phillips, E. L. Rall, and D. A. Peluso, "A Data-Frame Theory of Sensemaking," in *Expertise Out of Context: The Sixth International Conference on Naturalistic Decision Making*, R. R. Hoffman, Ed. Lawrence Erlbaum Associates, Inc., Sep 2007, pp. 13–155.

[12] G. A. Klein and R. R. Hoffman, "Seeing The Invisible: Perceptual–Cognitive Aspects of Expertise," in *Cognitive Science Foundations of Instruction*, M. Rabinowitz, Ed. Erlbaum, Oct 1992, pp. 203—-226. [Online]. Available: http://cmapspublic3.ihmc.us/rid=1G9NSY15K-N7MJMZ-LC5/SeeingTheInvisible.pdf

[13] G. Klein, *Sources of Power: How People Make Decisions*. MIT Press, 1998.

[14] G. Klein, R. Pliske, B. Crandall, and D. D. Woods, "Problem detection," *Cognition, Technology & Work*, vol. 7, no. 1, pp. 14–28, mar 2005. [Online]. Available: http://link.springer.com/10.1007/s10111-004-0166-y

[15] SANS, "Stop that phish," https://www.sans.org/security-awareness-training/resources/stop-phish, accessed: 2020-12-19.

[16] H. R. Lipford and M. E. Zurko, "Someone to watch over me," in *Proceedings of the 2012 New Security Paradigms Workshop*, ser. NSPW '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 6776. [Online]. Available: https://doi.org/10.1145/2413296.2413303

[17] P. DiGioia and P. Dourish, "Social navigation as a model for usable security," in *Proceedings of the 2005 Symposium on Usable Privacy and Security*, ser. SOUPS '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 101108. [Online]. Available: https://doi.org/10.1145/1073001.1073011

[18] J. Heurix, P. Zimmermann, T. Neubauer, and S. Fenz, "A taxonomy for privacy enhancing technologies," *Computers & Security*, vol. 53, pp. 1–17, 2015.

APPENDIX

A. Sample phishing email

[ACTION REQUIRED] Norbert Nthala: Meetings via MSU Zoom

○ MSU IT Services <noreply@itservices-msu.com>   Today at 12:46 PM
To: ✓ Nthala, Norbert

Dear Norbert,

MSU IT Services requests you to provide important information regarding your use of Zoom for work-related meetings at MSU. Due to increased use of Zoom for online teaching, most MSU staff and students are increasingly facing connection problems. MSU IT Services is urgently upgrading its Zoom subscription and making changes to balance administrative and academic use of the software. These changes will affect how and when you can schedule and attend meetings and they will be applied automatically to your MSU Zoom account on Friday, December 11. The changes will remain in effect until the university fully reopens.

To ensure your settings are customized according to your work requirements, you are required to provide MSU IT with the necessary information via this form https://forms.msu.edu/MCWw5NVv2TDgQUkd7 .

We appreciate the opportunity to serve you.

MSU Information Technology

Questions or concerns?

MSU IT is available 24/7 to support your IT needs. Contact the MSU IT Service Desk using one of the methods listed below.

· Call (517) 432-6200

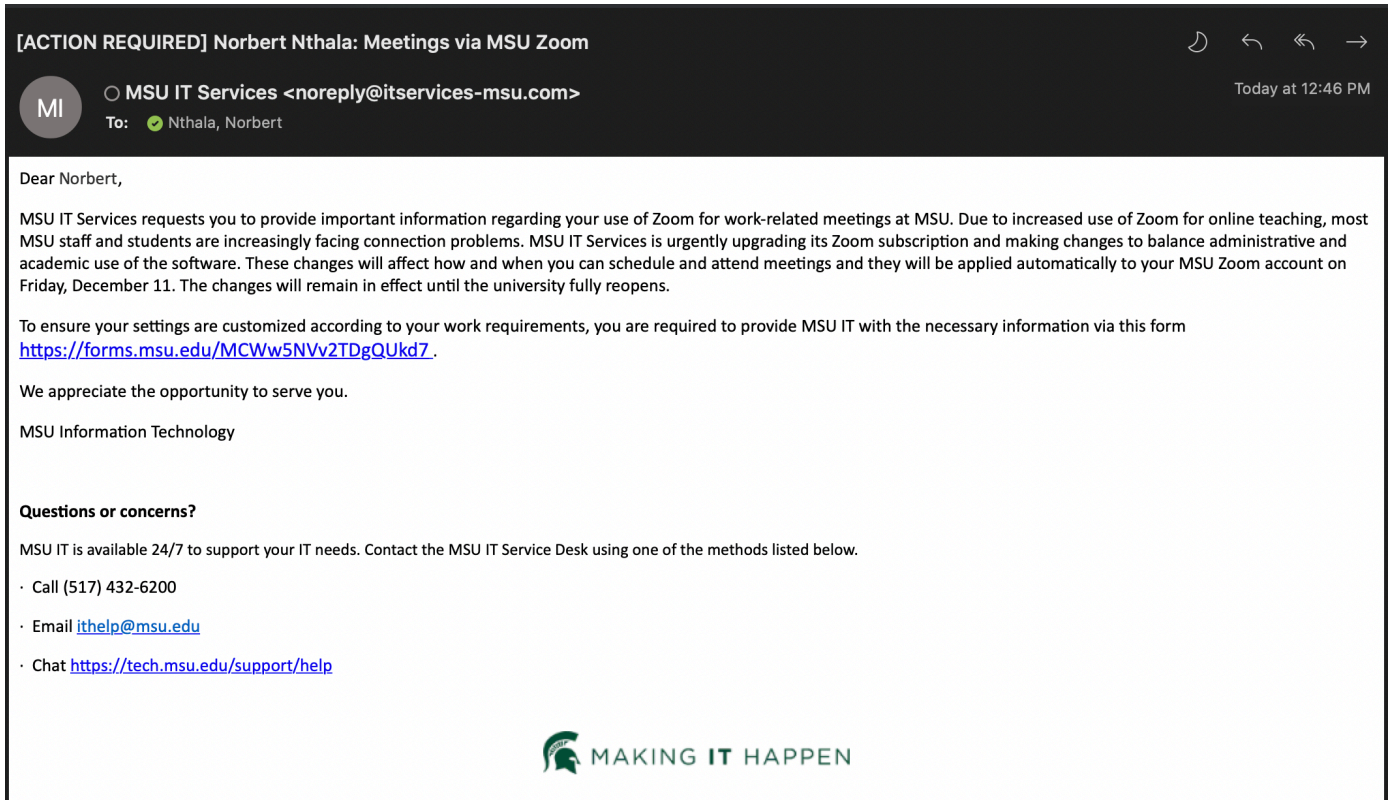· Email ithelp@msu.edu

· Chat https://tech.msu.edu/support/help

MAKING **IT** HAPPEN

Fig. 1. Sample phishing email