

Lessons From a Real World Evaluation of Anti-Phishing Training

ABSTRACT

Prior laboratory studies have shown that PhishGuru, an embedded training system, is an effective way to teach users to identify phishing scams. PhishGuru users are sent simulated phishing attacks and trained after they fall for the attacks. In this current study, we extend the PhishGuru methodology to train users about spear phishing and test it in a real world setting with employees of a Portuguese company. Our results demonstrate that the findings of PhishGuru laboratory studies do indeed hold up in a real world deployment. Specifically, the results from the field study showed that a large percentage of people who clicked on links in simulated emails proceeded to give some form of personal information to fake phishing websites, and that participants who received PhishGuru training were significantly less likely to fall for subsequent simulated phishing attacks one week later.

This paper presents some additional new findings. First, people trained with spear phishing training material did not make better decisions in identifying spear phishing emails compared to people trained with generic training material. Second, we observed that PhishGuru training could be effective in training other people in the organization who did not receive training messages directly from the system. Third, we also observed that employees in technical jobs were not different from employees with non-technical jobs in identifying phishing emails before and after the training. We conclude with some lessons that we learned in conducting the real world study.

Categories and Subject Descriptors

D.4.6 Security and protection, H.1.2 User / Machine systems, H.5.2 User interfaces, K.6.5 Security and protection education.

General Terms

Design, Experimentation, Security, Human factors.

Keywords

Embedded training, phishing, email, usable privacy and security, real world studies.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '04, Month 1–2, 2004, City, State, Country.
Copyright 2004 ACM 1-58113-000-0/00/0004...\$5.00.

1 INTRODUCTION

User education is a frequently-recommended and widely-used approach to countering phishing attacks [1, 12, 33], but few studies have evaluated the effectiveness of this approach in the real world. Researchers have demonstrated the effectiveness of PhishGuru, an embedded training system [20, 21]; and Anti-Phishing Phil, an online game [31] in laboratory studies. However, laboratory studies are unable to fully replicate real world conditions: they may lack ecological validity and do not sufficiently approximate real-world situations, which in turn may impact external validity - that is, the ability to make generalized inferences from the results [3]. The focus of this paper is to build on the earlier PhishGuru laboratory studies by conducting a similar study in a real world setting.

PhishGuru motivates users to pay attention to anti-phishing training materials by taking advantage of teachable moments. PhishGuru users are sent simulated phishing attacks via email and are presented training materials when they fall for the attacks. These emails might be sent by a corporate system administrator, ISP, or training company. The training materials present the following concepts in the form of a comic script: the definition of phishing, steps to follow to avoid falling for phishing attacks, and how criminals conduct phishing attacks easily.

Our goal is to evaluate the effectiveness of PhishGuru training in field trials and to study the effect of variations in the content of the PhishGuru training messages. To evaluate PhishGuru in the real world, we conducted a study with employees in a Portuguese company. The simulated phishing emails were all spear phishing emails targeted at the employees of the company. To investigate the effect of different training messages, we used one that had instructions on how to protect against regular phishing scams (generic training) and one that had instructions for protecting against spear phishing scams (spear training).

Our results demonstrate that the findings of PhishGuru laboratory studies do, indeed, hold up in the real world. As with the laboratory studies, our field study results showed that a large percentage of people who clicked on links in simulated emails proceeded to give some form of personal information to fake phishing websites, and that participants who received PhishGuru training were significantly less likely to fall for subsequent simulated phishing attacks one week later. In addition, we found the people trained with the spear phishing training material did not make better decisions in identifying spear phishing emails compared to people trained with the generic training material.

The remainder of the paper is organized as follows: In the next section we describe related work, including several training

methods, and some relevant experimental studies. In Section 3, we present the study setup, participant demographics, and hypotheses that guided our study. In Section 4, we present the results of our evaluation, demonstrating that PhishGuru is effective in educating people in the real world. We discuss the effect of training people in the real world in Section 5. In Section 6, we present some limitations along with lessons learned. Finally, we present our conclusions and future work in Section 7.

2 BACKGROUND

In this section we present an overview of security training methods; describe several methods for studying users' behavior in the context of phishing; and describe other experimental studies that have been conducted to evaluate the effectiveness of phishing training.

2.1 Security training methods

ISO and NIST security standards, which many companies are contractually obligated to follow, include security training as an important component of security compliance [13, 26]. These standards describe a three-level framework that includes awareness, training, and education. *Security awareness* activities are intended for all employees of a company and often include videos, newsletters, and posters. *Training* is generally intended only for employees who are involved with IT systems, mainly to provide basic computer security knowledge. Training is delivered primarily through classroom lectures, e-learning materials, and workshops. *Education*, intended for IT security specialists, is usually delivered via seminars or reading groups [25]. Our research offers some new approaches to delivering security awareness and training effectively.

There are many approaches to training users about phishing, including: articles about phishing on websites [8, 9, 10, 24], online cartoons about security [32], web-based phishing IQ tests [23], classroom training [28] and sending security notices in emails. These approaches vary in their cost as well as their effectiveness. For example, classroom training may be more effective than other training approaches because employees are required to spend dedicated time for training, but this approach is time-consuming for employees and expensive for companies that have to train a large number of employees. Online training materials are often an inexpensive approach, but it can be difficult to get people to read these materials and they are not always effective. PhishGuru approach is to present the training materials when people fall for phishing emails. This approach creates motivation among people and therefore can be effective in training.

2.2 User study methods

In order to develop effective anti-phishing training materials it is essential to understand why users fall for phishing attacks and how anti-phishing tools and training materials impact their behavior. Researchers have used a variety of methods in user studies designed to gain insights into these issues. Interview studies have been conducted to gain insights into users' mental models and decision processes [7, 18]. Laboratory experimental studies where participants played a fictitious role and used personal information associated with that role have been used to test users' susceptibility to phishing attacks and evaluate the effectiveness of anti-phishing toolbars and training materials [2, 6, 14, 19, 20, 21, 31]. Laboratory experimental studies where

participants used their own credentials have been used to evaluate the effectiveness of mutual authentication tools [30]. Real world studies have been used to evaluate participants' susceptibility to phishing, but not to evaluate the effectiveness of training [11, 15, 27].

Laboratory experimental studies are very helpful in understanding the user behavior in a given situation. However, each of these study methods have tradeoffs and face validity challenges: most of these studies are challenged with ecological (whether the methods, materials, and settings are similar to real life) and external (whether the results are generalizable) validity issues [3]. Laboratory experimental studies in the context of phishing are also challenged with ethical issues of how much the researcher should inform the participant about the study and how much deception is acceptable [16, 17]. In one laboratory experimental setup, researchers showed that people who role-play behave differently from people who use their own credentials [30].

Understanding users' behavior in real world settings is critical to developing effective counter measures for phishing. Even though real world studies provide richer data, it may be difficult to control the study setup (due to many sources of variability) in the real world [29]. It can also be difficult to make the arrangements for a real world study, especially when it requires the cooperation of a company to gain access to employees or customers. Companies may not grant desired access or permit publication of study data or results.

2.3 Experimental evaluation of anti-phishing training

Few real world studies of users' behavior in the context of phishing have been conducted, and even fewer real world studies have been conducted to evaluate the effectiveness of anti-phishing training. Real world evaluations of anti-phishing training involve classroom and office training as well as training delivered via an online game. Researchers have evaluated the effectiveness of security notices and embedded training in laboratory studies.

The idea of sending fake phishing emails to test users' vulnerability has been explored by several groups. Jagatic et al. conducted a study in which they sent phishing emails to Indiana University students that pretended to come from one of their friends (this information was obtained through social networking websites). A large percentage of students fell for these phishing attacks [15]. Ferguson did a two-part study among West Point cadets. In the first phase, cadets were tested for their ability to detect phishing attacks. In the second phase, cadets were given classroom training and lectures about phishing and then tested. Ferguson showed an improvement in the cadets' ability to identify phishing emails after the training [11]. Similar to the West Point cadet study, the New York state office of Cyber Security & Critical Infrastructure Coordination conducted a two-part study among their employees. In this study, participants who fell for simulated phishing attacks were presented with online educational materials on how to protect themselves from phishing. This study also showed anti-phishing training improved participants' ability to identify phishing emails [27].

Sheng et al. have shown that people can be trained about phishing URLs through an online game called Anti-Phishing Phil. In a laboratory study, they found that users made better decisions when trained with the game than with existing online training

materials [31]. They found similar results while testing the game in the real world [22].

Previous research results provide strong evidence that people make better decisions when they are trained through embedded training versus the current practice of sending security notices [20]. Research also suggests that people retain and transfer more knowledge when trained with embedded training than with non-embedded training [21]. The focus of this paper is on testing embedded training in the real world.

3 EVALUATION

In this section we present participant demographics and study methodology along with the hypotheses that we tested in this study.

3.1 Participants and demographics

This study was conducted at a Portuguese company. All emails and the training materials were translated into Portuguese. All participants in the study worked in the same floor of an office building. Participants were from different areas of work in the company: administration, commercial, design, editorial, management, technical, and others.

The study included three conditions: “control,” “generic training,” and “spear training.” Participants in the control condition did not receive any training. Participants in the generic training condition received a simulated spear phishing email and saw generic phish training material (Figure 1) when they clicked on a link in the email. Participants in the spear training condition received a simulated spear phishing email and saw spear phish training material (Figure 2) when they clicked on a link in the email. We assigned 111 employees to the control condition, 100 to the generic training condition, and 100 to the spear training condition. Table 1 presents the demographics of the study participants.

3.2 Study setup

The company we worked with was primarily interested in studying the vulnerability of their employees towards spear phishing emails, so we used spear phishing emails for all simulated phishing emails in this study. Targeted spear phishing attacks have been more successful than generic phishing attacks in coning people and causing damages to companies and individuals.

In total, participants received four emails during the study: three simulated spear phishing emails and one legitimate email containing a link. All the spear phishing emails and the legitimate email were based on actual emails that the company had received

Table 1: Demographics of the participants.

	Control Condition (N=111)	Generic training condition (N=100)	Spear training condition (N=100)
Gender			
Male	77%	27%	67%
Female	23%	73%	33%
Areas of work			
Administration	1%	1%	1%
Commercial	2.7%	5%	9%
Design	5.4%	3%	7%
Editorial	4.5%	5%	7%
Management	22.5%	19%	20%
Technical	39.6%	36%	35%
Others	24.3%	31%	21%

or the kind of emails that the system administrators were worried about.

The first email that employees received was a training email (Train) and was delivered on Day 0. This email was sent only to employees in the generic and spear conditions. This email was a spear phishing email that asked employees to click on a link to enter their user name and password in order to use the corporate network. When employees clicked on the link in this email, they were taken to the training material corresponding to the condition they were in. Participants in the generic training condition saw the generic phish training message shown in Figure 1, while participants in the spear training condition saw the spear phish training message shown in Figure 2.

The second email (Test 1) was designed to measure the knowledge that employees acquired through our training materials. In order to compare trained and untrained employees, this email was sent to employees in all conditions. To measure immediate effectiveness this email was sent on Day 2 of the study. This simulated spear phishing email said that the recipient’s internal network password has expired and asked them to click on a link and change their password. When employees clicked on link in this email, they were taken to a fake phishing website that looked the same as the real website and was hosted in a similar looking domain name.

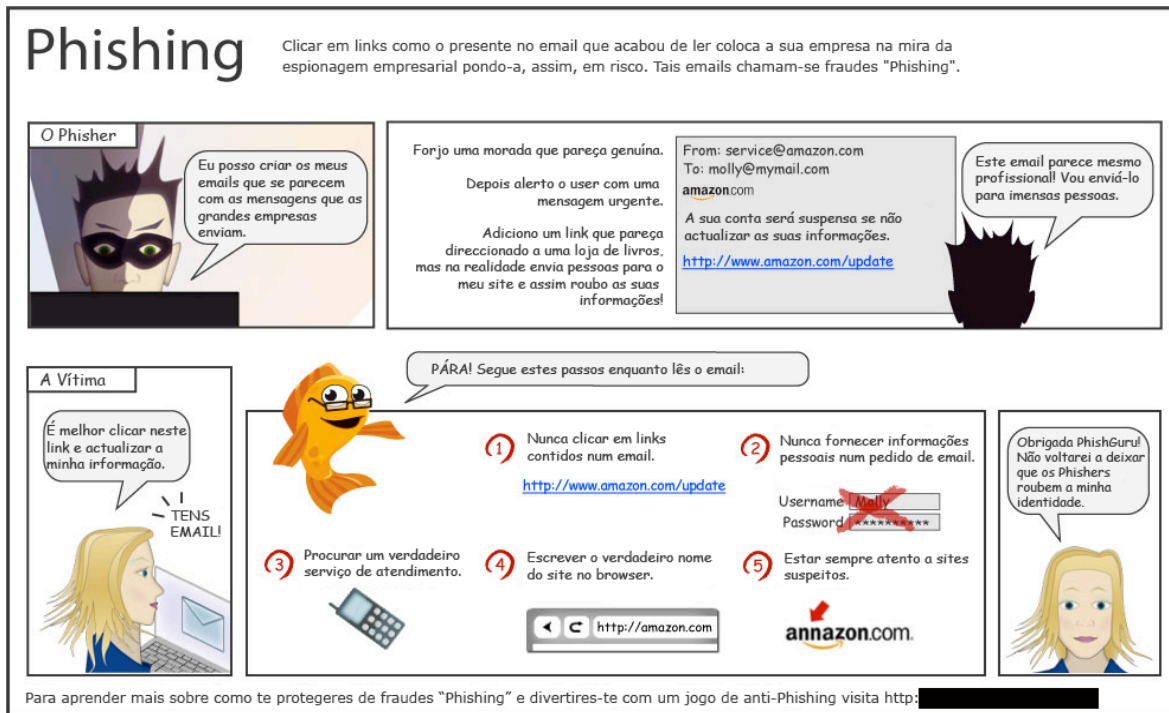


Figure 1. People in Generic condition saw this comic strip.

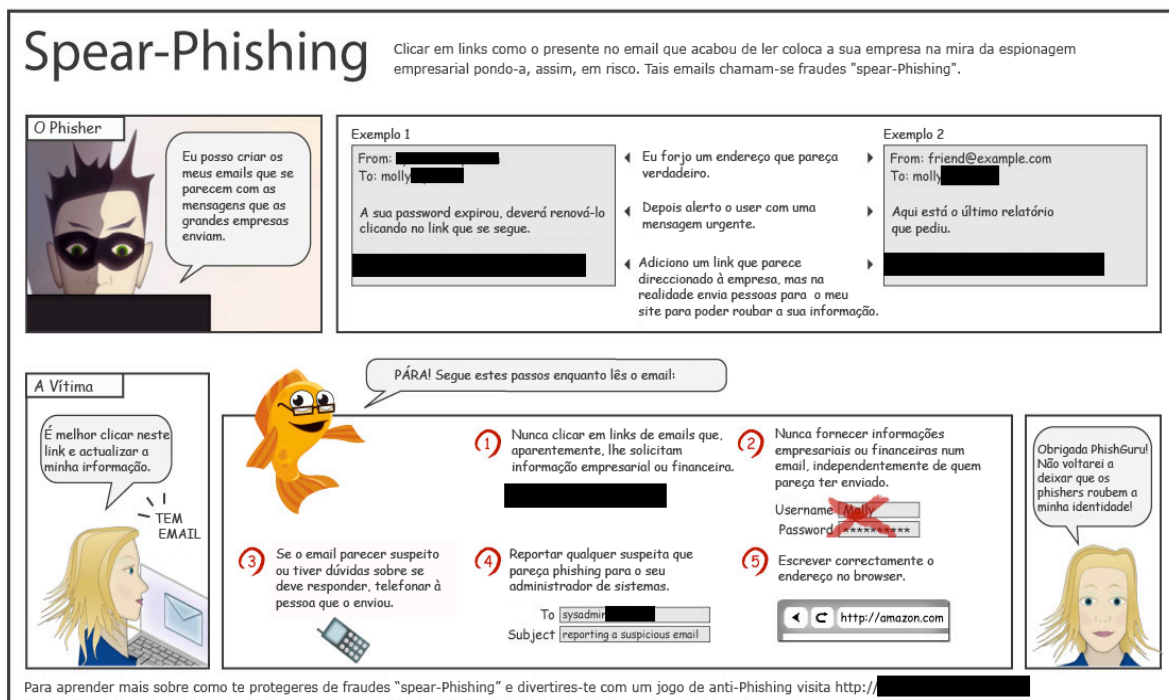


Figure 2. People in Spear condition saw this comic strip

Learning science literature defines retention as the ability of learners to retain or recall the concepts and procedures taught when tested under the same or similar situations after a time

period δ from the time of knowledge acquisition [5]. The third email (Test 2), which was designed to measure retention, was sent on Day 7. As in Test 1, to compare the trained and untrained

Table 2: Summary of emails sent to study participants

Emails	Type	Day of sending	Conditions	Relevant features of the email
Train	Spear phishing	Day 0	Generic and spear	To enter their user name and password in order to use the corporate network
Test 1	Spear phishing	Day 2	All	Internal network password expired; to change their password
Test 2	Spear phishing	Day 7	All	To update their communication information
Test 3	Legitimate with link	Day 10	All	To read the updated security policy of the company

employees, this email was sent to participants in all conditions. This email asked employees to click on a link and update their communication information for internal corporate communication purposes. When employees clicked on the link they were taken to a phishing website that looked the same as the real website and was hosted on a similar looking domain name.

To test whether training increases participants' concern level such that they stop clicking on any links, even legitimate ones, we sent a legitimate email with a link (Test 3) on Day 10 to all participants in the study. To compare the trained and untrained employees, this email was sent to participants in all conditions. This email asked employees to click on a link to read the company's updated security policy. When employees clicked on the link, they were taken to a legitimate webpage with the updated security policy. Table 2 summarizes all emails, email types, days on which the email was sent, the conditions to which the emails were delivered, and relevant features of the email.

Phishing websites that were linked to the spear phishing emails were exact replica of real company websites but hosted on a domain that looked similar but not the same as the company's domain. All replicated websites were completely functional and allowed employees to submit information. We wanted only the employees of the company to access the training materials and fake phishing websites, so, these websites were hosted in a way that only IP addresses coming from the company's domain were able to access the websites. This also helped us in identifying the IP address and thereby the user from whose machine the request had come. We tracked the clicks to the phishing websites and the training materials, as well as the information that was submitted to the phishing websites.

To make sure the employees received the emails that were part of the study, system administrators bypassed the corporate email filters and placed them in participants' inboxes.

We asked all participants to complete a post-study survey on Day 20. The survey consisted of questions regarding (1) the interest level of participants in receiving such emails in future, (2) participants' feedback on the training, and (3) participants' feedback on the instructions.

3.3 Hypotheses

In this section we introduce three hypotheses which informed the study described in the paper.

3.3.1 Replicating laboratory study results

Earlier laboratory studies have shown that a large percentage of participants who click on links in simulated emails proceed to give some form of personal information to the phishing website. This percentage was around 90% in earlier laboratory studies [20, 21]. Our goal was to investigate whether this is true in a real world setting. This result may show that people have to be trained not to click on links, otherwise, there is low probability that they will click and not give personal information to phishing websites.

Hypothesis 1: *A large percentage of people who click on links in simulated emails proceed to give some form of personal information in the real world.*

A laboratory study showed that users learn, retain, and transfer effectively when training materials are presented after they fall for a phishing attack [21]. Our goal was to investigate whether this is true in a real world setting.

Hypothesis 2: *PhishGuru (embedded training) is effective in*

Table 3: Translated English version of the instructions in the training materials.

Generic training instructions	Spear training instructions
1. Never click on links within emails	1. Never click on links within emails that appear to be requesting corporate or financial information
2. Never give out personal information upon email request	2. Never give your corporate or financial information over the email, no matter who appears to have sent it
3. Find and call a real customer service center	3. If an email looks suspicious or you are uncertain about whether to respond, call the person who sent it
4. Type in the real website address into a web browser	4. Report any suspicious email that could be spear phishing to sysadmin@company.com
5. Always be wary of suspicious websites	5. Type in the real website address into a web browser

training people in the real world.

To evaluate the effectiveness of PhishGuru, we calculated the following: (1) percentage of participants who clicked on a link in phishing emails and gave information to fake phishing websites immediately after the training; (2) percentage of participants who clicked on a link in phishing emails and giving information to fake phishing websites after a delay of 7 days from the training; and (3) percentage of participants who clicked on a link in legitimate emails after the training.

3.3.2 Generic and spear training instructions

The content of training materials makes a difference in the way people learn and reproduce knowledge. Researchers have shown that people make better decisions if the testing situation is the same or similar to the training situation and the training materials than if the testing situation is different [5]. To investigate the effect of the difference in the instructions, we developed one set with anti-phishing instructions that were generic and another one specific to spear phishing emails. Figure 1 and Figure 2 have the same content except for the instructions in the lower pane of the material. As the training materials used in the study were in Portuguese, the translated English version of the instructions is given in Table 3.

Hypothesis 3: *People trained with spear training material make better decisions in identifying spear phishing emails compared to people trained with generic training material.*

4 Results

In this section we present the results of our study. The results from this study support Hypotheses 1 and 2, but not Hypothesis 3. We found a large percentage of the participants who clicked on links in simulated emails gave away some form of personal information to the fake phishing websites that were part of the study. We found participants in the training conditions made significantly better decisions after the training compared to before the training. Our results suggest that users retained knowledge gained from PhishGuru for at least 7 days after the training. However, the difference in the instructions in our training materials did not have a significant effect on the participants' ability to identify phishing emails. Surprisingly, our results also suggest that PhishGuru training could be effective in training other people in the organization who did not receive training messages directly from the system. The complete decision tree for all the three conditions is given in the Appendix.

4.1 Giving away personal information

In this study we found that a large percentage of the participants who clicked on links in simulated phishing emails went ahead and gave some form of personal information to the phishing websites. The system administrators in the company who helped us conduct the study had access to the information that was entered into phishing websites. They were able to check the usernames and other details that were entered. We found that 88% of the participants who clicked on links went ahead and gave some form of personal information to the fake phishing websites. In laboratory studies, researchers have found that 90 to 93 percent of participants who clicked on links gave their personal information to fake phishing websites [20, 21]. Table 4 gives the percentage of participants in each condition who clicked on a link in phishing emails, and who clicked and gave information to fake phishing websites.

4.2 Phishing emails

We found that PhishGuru training improved participants' decision making on the phishing emails that they received as part of the study. Before training, we see (Table 4) no significant difference between generic (42%) and spear (39%) training conditions for the percentage of participants who clicked on link in the phishing email and gave information (two sample T-test, p-value = 0.6). This shows that before the training, participants were at the same level in both conditions.

In both the training conditions (generic and spear), participants acquired and made improved decisions immediately after training. We found (Table 4), in the generic condition, the percentage of participants clicking and giving information reduced significantly from 42% on Day 0 to 15% on Day 2 (paired T-test, p-value < 0.01). In the spear training condition the percentage decreased significantly from 39% on Day 0 to 12% on Day 2 (paired T-test, p-value < 0.01).

Trained participants (who clicked on the link in Train email and saw the training materials) retained the knowledge gained from PhishGuru training for at least 7 days after the training. Table 5 gives the percentage of those participants who got trained and who clicked on link and gave information. The untrained group includes participants both from generic training and spear training conditions who did not see the training materials. From Table 5, we see that participants did not lose significant knowledge on Day 7 compared to Day 2 in the generic training condition (Paired T-test, p-value = 0.55) or in the spear training condition (Paired T-test, p-value = 0.67).

Table 4: Percentage of participants who clicked on the training link, only clicked, and clicked and gave information on other emails.

Conditions	Clicked on link in training email on Day 0	Clicked on link on Day 2	Clicked on link and gave information on Day 2	Clicked on link on Day 7	Clicked on link and gave information on Day 7
Control	N/A	20 %	19 %	17 %	15 %
Generic training	42 %	17 %	15 %	14 %	12 %
Spear training	39 %	14 %	12 %	17 %	14 %

Table 5: Percentage of those participants who clicked on link on Day 0, and clicked on link and gave information on Day 2 and Day 7.

	Day 0	Day 2	Day 7
Generic trained	100 %	19 %	12 %
Spear trained	100 %	18 %	15 %
Untrained	0 %	10 %	13 %

Table 6: Percentage of participants correctly (did not click on the link in the email) identifying the Day 2 and Day 7 emails. The untrained group includes participants from both training groups who did not actually receive training.

Conditions	Identified 2 emails correctly	Identified 1 email correctly	Identified 0 email correctly
Control	58.2 %	32.8 %	8.9 %
Generic trained	70.4 %	18.5 %	11.1 %
Spear trained	65.2 %	30.4 %	4.3 %
Untrained	73.4 %	22.8 %	3.8 %

We found that a significant number of trained participants identified both of the test emails correctly. Table 6 shows the percentage of control, trained, and untrained participants who identified Day 2 and Day 7 emails correctly. The untrained group includes participants from both the generic and spear training conditions who did not see the training materials because they did not click on the link in the first phishing email. In the trained conditions, we see significant number of participants identified both emails correctly. We believe that retraining with a second training email could further improve the percentage of participants who could identify both emails correctly. Our results also show that untrained participants identified phishing emails better than trained participants. This suggests that most of these participants did not need the training that they did not receive.

These results demonstrate that participants in the generic and spear training conditions were able to make improved decisions immediately after being trained and they were able to retain the knowledge for at least 7 days.

4.3 Legitimate emails

We do not have enough data to conclude whether training increased the concern level of the participants so much that they refrained from clicking on any email links, even legitimate ones. We found only three employees across all the three conditions who clicked on the link in the legitimate email that was sent as part of the study on Day 10. In order to verify this behavior, we sent another legitimate email on Day 14 from the marketing team, with a link to a company sales report. Again, only three employees across all conditions clicked on the link in the legitimate email. There was no difference between control and training (generic and spear) conditions. This suggests that the behavior we observed may not be the effect of training, but rather

the normal behavior of employees in this company towards such corporate emails.

The content of the training and testing emails used in the study has to be properly designed to provide incentives for the participants. Employees in the company may not read email messages unless they are very relevant to them or has severe consequences. Ideally it would have been useful if we could have sent a legitimate email before the training to understand the baseline. Since we don't have the baseline data of how participants respond to their legitimate emails, we will not be able to use this to either support or reject Hypothesis 2.

4.4 Generic vs. spear instructions

Our results suggest that the difference in the instructions that we had in our training materials did not have an effect on the participants' ability to identify phishing emails. From Table 4, we see that percentage of participants who clicked the link and gave information on Day 2 for the generic training condition was not significantly different from the spear training condition (two sample T-test, p-value = 0.53). Similarly, we found the difference on Day 7 also to be insignificant (two sample T-test, p-value = 0.67). In Table 5 we examine only those participants in the generic and spear conditions who actually received training. We see that there was no significant difference between the trained conditions for the test email on Day 2 (two sample T-test, p-value = 0.8) or Day 7 (two sample T-test, p-value = 0.7). This suggests that participants don't gain specific ability for identifying phishing emails by seeing specific instructions rather than generic instructions.

Using both the total percentage (Table 4) and the percentage of employees who got trained (Table 5), we found no significant difference between employees in generic and spear condition in their ability to identify phishing emails. Thus we must reject Hypothesis 3. However, we believe this hypothesis warrants further investigation. A more substantial difference between the generic and spear training might produce a significant effect. In addition, because all of the participants in this study worked on the same floor of an office building, we are concerned that participants across conditions may have shared the training materials they received with each other. Further investigation is needed to understanding the influence of instructions on decision making.

4.5 Observations

We have anecdotal evidence that employees discussed the study among themselves and with their system administrators, and we believe this had an impact on our results. Although only 50 employees clicked on the training material link, our logs show that the material was downloaded 95 times during the study (which means that some employees viewed the training material multiple times). Some people may have shown the training to colleagues in other conditions. We believe this is likely to have caused participants in the control condition make right decisions on Day 2 and Day 7, even though they received no direct training. However, they may have received indirect training when participants in the training conditions told them about their training or showed them the training messages. We have anecdotal evidence that employees did not receive any other information about phishing from the company and there was no drastic incident that could have influenced employees to change their behavior during the study. This suggests that PhishGuru training can be effective in training people who are not part of the

study – it may be good enough to train a subset of employees who may influence other employees in the company. Researchers have shown that physical proximity and social structure of people may trigger information flow [4]. We attribute this result to the way the employees were seated in the company – all employees were on the same floor. Further investigation may explain this phenomenon better.

Job type did not have any influence on participants' ability to identify phishing emails either before or after the training. In particular, we compared technical and non-technical job types. Before the training, the percentage of participants in the generic training condition who clicked the link and gave information was the same (42%) for technical and non-technical employees. For the spear condition, this percentage is 48% for technical and 34% for non-technical participants. This difference was statistically insignificant (two sample T-test, p -value = 0.16). Similarly, we found no significant difference between technical and non-technical employees after the training.

We found no significant difference in susceptibility to phishing emails between male and female employees (Two sample T-test, p -value = 0.76). Other researchers have found similar results [6, 21, 31].

We circulated a post-study questionnaire to participants to get their feedback about PhishGuru training and the training materials. Unfortunately none of the employees turned in their completed questionnaire. In future studies, we plan to give some incentive for the participants to fill out the post-study questionnaire.

5 Discussion

The results from the study supported Hypothesis 1 and Hypothesis 2 (for phishing emails and needs further investigation for the legitimate emails). Further research is needed to investigate Hypothesis 3.

Our results are consistent with earlier laboratory studies that demonstrated the effectiveness of the PhishGuru embedded training system. Our results suggest that a large percentage of people who clicked on links in emails proceeded to give some form of personal information. Other researchers have found the same in laboratory studies [21]. Our results also strongly suggest that PhishGuru is effective in training employees in the real world. In the earlier studies researchers have shown that users were motivated to learn when the training materials are presented after users fall for the phishing emails (when users click on the link in the email) compared to sending instructional materials through email (non-embedded). In this paper we showed users' ability to identify phishing emails improved after the training. Due to lack of data we were not able to conclude anything about legitimate emails, therefore, there needs further investigation on whether training increases the concern level of the participants in the real world. Our results also suggest that employees retained for at least 7 days the knowledge that they gained by reading the training material. Other researchers have also found similar results in laboratory studies [21]. Our results showed that significant number of participants identified both the testing emails correctly compared to participants identifying one or none correctly.

A laboratory study [21] showed that 79% of the participants clicked and gave information before the training while this was 41% in the real world. Seven days after the training, the percentage reduced to 35% in lab study and 13% in the real world.

This observed differences between the laboratory and real world studies may be due to differences in demographics of the participants, difference in language of the study materials (English versus Portuguese), or differences in the simulated phishing emails used. It may also be due to the fact that real world participants are using their own credentials while in the lab they use fictitious details. Despite the initial differences, participants in both the laboratory and real world study showed similar abilities to learn from the training materials.

Our results suggest that there is no significant difference between employees who got trained through generic training instruction and spear training instruction in identifying phishing emails. We believe that this may be due to small sample size of employees who were trained and also who clicked on the link and gave information for the testing emails, and employees discussing among themselves about the study. We hope to investigate the effect of difference in training materials in future studies.

The results also showed that employees discussed the phishing emails and the training materials that were sent as part of this study among themselves. This may not be good for our study, but it suggests that by training a subset of employees, a company can expect these trained employees to influence other employees who were not part of the training. It would have been useful if we had more data to show this effect, but this may be a good starting point for further investigation on this topic.

6 Limitations and lessons learned

As in laboratory studies, our field study also had a few limitations. We did not send a simulated phishing email to the control condition on Day 0, so, we do not have baseline data for all participants before training. If we had this data we could have measured the effect of training more directly by comparing participants in different conditions who clicked the link in the Day 0 email. In future studies we plan to send a simulated phishing email linked to a functional fake phishing website on Day 0 to the control condition. We also did not send a legitimate email before the training and therefore we could not understand the behavior of the participants towards legitimate emails before the training. This restricts us from explaining the low percentage of participants who clicked on the link in the legitimate email.

There were many lessons learned which will help future studies:

- *Content of the email is important:* The simulated emails that are used in the study should be relevant and have a compelling argument for participants to make a decision.
- *Incentive for participants:* The employees who are part of the study may not have the incentive to provide feedback or complete an exit survey. So, providing some form of incentive (cash or prize) to the participants is necessary.
- *Use participants from different locations:* It is useful to select study participants from different work locations so they are less likely to discuss the study among themselves.
- *Keep it simple:* The companies that agree to do real world studies may not have the incentive to collect data at the level of detail that researchers would want. Therefore, the procedures for collecting data should be minimized and made simple.

7 Conclusion and future work

In this paper we presented the first empirical evaluation of embedded training methods that teach people about phishing

during their normal use of email in the real world. In this paper we showed that: (a) a large percentage of people who click on links in simulated emails proceed to give some form of personal information in the real world; (b) PhishGuru training, an embedded training, is effective in training people in the real world; (c) users retained knowledge for at least one week when trained with embedded training in the real world; (d) people trained with spear training instruction did not make better decisions in identifying spear phishing emails compared to people trained with generic training instruction.

Based on lessons that we learned from this study, we are currently designing a study with another company. We plan to collect richer data. We are also currently working on designing instructional materials for other cues and strategies that users can be trained on.¹

8 ACKNOWLEDGMENTS

9 REFERENCES

- Allen, M. 1993. Social engineering: A means to violate a computer system. Tech. rep., SANS Institute, 2006.
- Anandpara, V., Dingman, A., Jakobsson, M., Liu, D., and Roinestad, H. 2007. Phishing IQ tests measure fear, not ability. Usable Security (USEC'07). <http://usablesecurity.org/papers/anandpara.pdf>.
- Brewer, M. Research Design and Issues of Validity. In *Handbook of Research Methods in Social and Personality Psychology*, pages 3 – 16. Cambridge University Press, 2000.
- Burt, R. S. 1987. Social contagion and innovation: Cohesion versus structural equivalence. *The American Journal of Sociology*, 92, 6, 1287–1335.
- Clark, R. C., and Mayer, R. E. 2002. E-Learning and the science of instruction: proven guidelines for consumers and designers of multimedia learning. John Wiley & Sons, Inc., USA.
- Dhamija, R., Tygar, J. D., and Hearst, M. 2006. Why phishing works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM Press, New York, NY, 581-590. DOI= <http://doi.acm.org/10.1145/1124772.1124861>.
- Downs, J. S., Holbrook, M. B., and Cranor, L. F. 2006. Decision strategies and susceptibility to phishing. In *Proceedings of the Second Symposium on Usable Privacy and Security* (Pittsburgh, Pennsylvania, July 12 - 14, 2006). SOUPS '06, vol. 149. ACM Press, New York, NY, 79-90. DOI= <http://doi.acm.org/10.1145/1143120.1143131>.
- eBay. Spoof Email Tutorial. Retrieved December 30, 2006. <http://pages.ebay.com/education/spoofutorial/>
- Federal Trade Commission. An E-Card for You game. Retrieved December 30, 2006. <http://www.ftc.gov/bcp/online/ecards/phishing/index.html>.
- Federal Trade Commission. Phishing Alerts. Retrieved December 30, 2006. <http://www.ftc.gov/bcp/online/pubs/alerts/phishingalrt.htm>
- Ferguson, A. J. 2005. Fostering E-Mail Security Awareness: The West Point Carronade. EDUCASE Quarterly, 1. Retrieved March 22, 2006, <http://www.educause.edu/ir/library/pdf/eqm0517.pdf>.
- Hiner, J. 2002. Change your company's culture to combat social engineering attacks. Retrieved Nov 3, 2006. http://articles.techrepublic.com.com/5100-1035_11-1047991.html.
- ISO. 2005. ISO/IEC 27001:2005 - Information technology – Security techniques – Information security management Systems – Requirements. Tech. rep., International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC), October 2005.
- Jackson, C., Simon, D., Tan, D., and Barth, A. 2007. An evaluation of extended validation and picture-in-picture phishing attacks. In Usable Security (USEC'07). <http://usablesecurity.org/papers/jackson.pdf>.
- Jagatic, T., Johnson, N., Jakobsson, M., and Menczer, F. Social phishing. 2007. Communications of the ACM 50, 10, 94–100. Retrieved March 7, 2006, <http://www.indiana.edu/phishing/social-network-experiment/phishing-preprint.pdf>.
- Jakobsson, M. and Ratkiewicz, J. 2006. Designing ethical phishing experiments: a study of (ROT13) rOnl query features. In *Proceedings of the 15th international Conference on World Wide Web* (Edinburgh, Scotland, May 23 - 26, 2006). WWW '06. ACM Press, New York, NY, 513-522. DOI= <http://doi.acm.org/10.1145/1135777.1135853>
- Karlof, C., Tygar, J., and Wagner, D. 2008. A user study design for comparing the security of registration protocols. Usability, Psychology, and Security (UPSEC).
- Kumaraguru, P., Acquisti, A., and Cranor, L. 2006. Trust modeling for online transactions: A phishing scenario. In Privacy Security Trust. http://www.cs.cmu.edu/onguru/pk_aa_lc_pst_2006.pdf.
- Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., and Hong, J. 2007. Teaching johnny not to fall for phish. Tech. rep., Cranegie Mellon University. <http://www.cylab.cmu.edu/files/cmucylab07003.pdf>.
- Kumaraguru, P., Y. Rhee, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. 2007. In *Proceedings of CHI 2007*. Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System.
- Kumaraguru, P., Rhee, Y., Sheng, S., Hasan, S., Acquisti, A., Cranor, L. F., and Hong, J. 2007. Getting users to pay attention to anti-phishing education: Evaluation of retention

- and transfer. e-Crime Researchers Summit, Anti-Phishing Working Group.
22. Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., and Hong, J. 2008. Under submission.
 23. Mail Frontier. Phishing IQ. Retrieved Sept 20, 2006. <http://survey.mailfrontier.com/survey/quiztest.html>.
 24. Microsoft. Consumer Awareness Page on Phishing. Retrieved September 10, 2006. <http://www.microsoft.com/athome/security/email/phishing.mspx>.
 25. NIST. 1998. Information technology security training requirements: A role- and performance-based model (800-16). Tech. rep., National Institute of Standards and Technology.
 26. NIST. 2004. NIST special publication 800-12: An introduction to computer security - the NIST handbook. Tech. rep., National Institute of Standards and Technology.
 27. New York State Office of Cyber Security & Critical Infrastructure Coordination. 2005. Gone phishing... a briefing on the anti-phishing exercise initiative for New York state government. Aggregate Exercise Results for public release.
 28. Robila, S. A., J. James and W. Ragucci. 2006. Don't be a phish: steps in user education. ITICSE '06: *Proceedings of the 11th annual SIGCSE conference on Innovation and technology in computer science education*. pp 237-241. New York, NY, USA.
 29. Salden, R., Aleven, V., Renkl, A., and Schwonke, R. 2008. Worked examples and tutored problem solving: redundant or synergistic forms of support? In Annual Meeting of the Cognitive Science Society. In press.
 30. Schechter, S. E., Dhamija, R., Ozment, A., and Fischer, I. 2007. The emperor's new security indicators. In SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy (Washington, DC, USA), IEEE Computer Society, pp. 51-65.
 31. Sheng, S., Magnien, B., Kumaraguru, P., Acquisti, A., Cranor, L. F., Hong, J., and Nunge, E. 2007. Anti-phishing Phil: The design and evaluation of a game that teaches people not to fall for phish. In SOUPS '07: *Proceedings of the 3rd symposium on Usable privacy and security* (New York, NY, USA, March 2007), ACM, pp. 88-99. Symposium On Usable Privacy and Security. DOI=<http://doi.acm.org/10.1145/1280680.1280692>.
 32. Srikwan, S., and Jakobsson, M. Using cartoons to teach security. Tech. rep., DIMACS, 2007. Retrieved April 1, 2008, <http://www.informatics.indiana.edu/markus/documents/security-education.pdf>.
 33. Timko, D. 2008. The social engineering threat. Information Systems Security Association Journal.

10 APPENDIX

Figure 3: Decision tree for control condition. Fallen refers to employees who clicked on link in the email and gave information and Not fallen refers to employees not clicking on links.

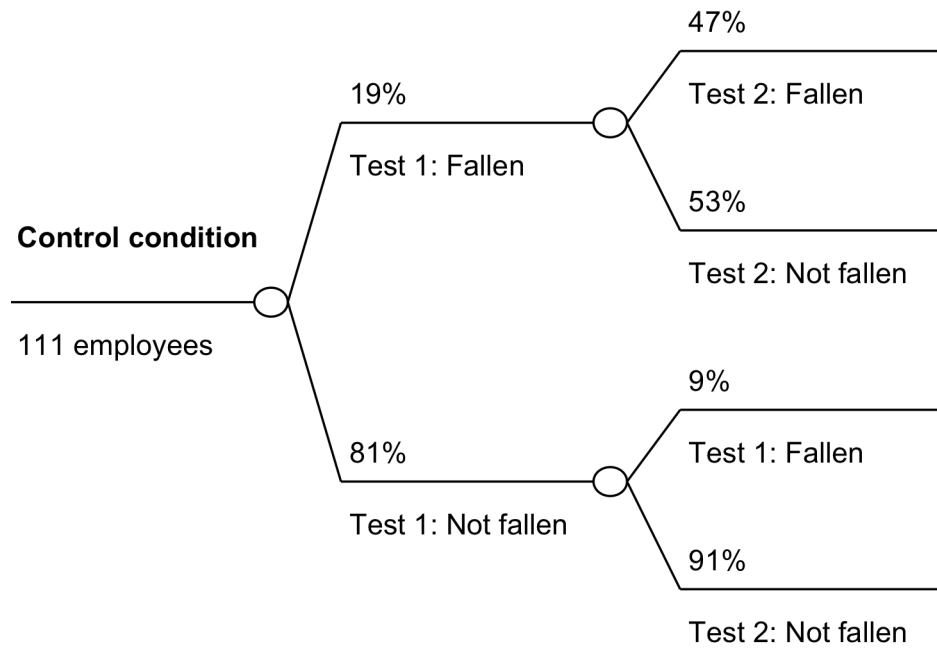


Figure 4: Decision tree for generic condition. Fallen refers to employees who clicked on link in the email and gave information and Not fallen refers to employees not clicking on links.

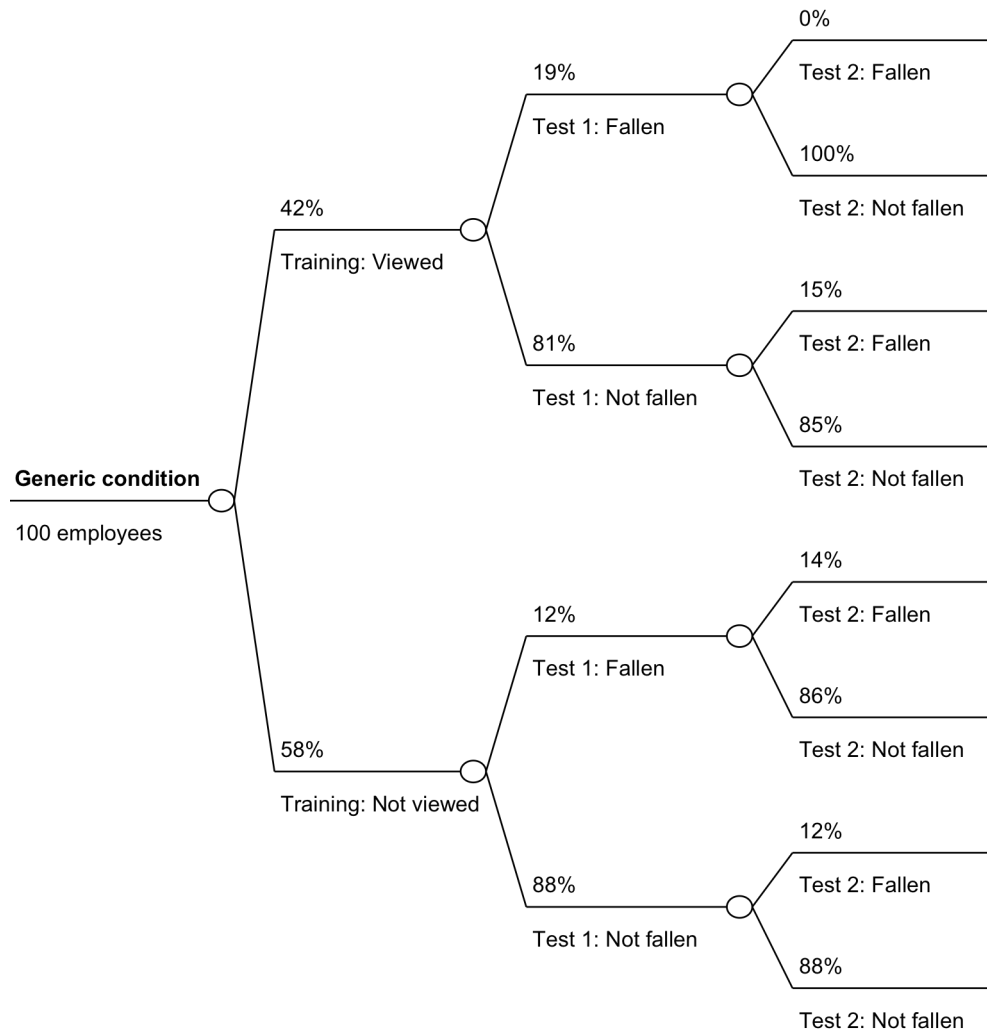


Figure 5: Decision tree for spear condition. Fallen refers to employees who clicked on link in the email and gave information and Not fallen refers to employees not clicking on links.

