Using Phishing Experiments and Scenario-based Surveys to Understand Security Behaviours in Practice

W.R. Flores[1], H. Holm[1], Gu. Svensson[1] and G. Ericsson[2]

[1]Department of Industrial Information and Control Systems, Royal Institute of Technology, Stockholm, Sweden
[2]Swedish National Grid, Stockholm, Sweden
e-mail: waldorf@ics.kth.se

Abstract

Threats from social engineering can cause organisations severe damage if they are not considered and managed. In order to understand how to manage those threats, it is important to examine reasons why organisational employees fall victim to social engineering. In this paper, the objective is to understand security behaviours in practice by investigating factors that may cause an individual to comply with a request posed by a perpetrator. In order to attain this objective, we collect data through a scenario-based survey and conduct phishing experiments in three organisations. The results from the experiment reveal that the degree of target information in an attack increases the likelihood that an organisational employee fall victim to an actual attack. Further, an individual's trust and risk behaviour significantly affects the actual behaviour during the phishing experiment. Computer experience at work, helpfulness and gender (females tend to be less susceptible to a generic attack than men), has a significant correlation with behaviour reported by respondents in the scenario-based survey. No correlation between the performance in the scenario-based survey and experiment was found. We argue that the result does not imply that one or the other method should be ruled out as they have both advantages and disadvantages which should be considered in the context of collecting data in the critical domain of information security. Discussions of the findings, implications and recommendations for future research are further provided.

Keywords

Social engineering, phishing, security behaviours, survey method, experiment

1. Introduction

The increased effectiveness and robustness of technical security components has made it more difficult to successfully attack computer systems using purely technical means. Many attackers have therefore started to include social means in their malicious efforts and target the humans accessing and using the computers (Applegate 2009). These types of attacks are commonly known as social engineering attacks. Social engineering is a form of deception in which an attacker attempts to deceive a victim into performing an action that benefits the attacker, e.g., click on a malicious link and install malware on their computers or reveal personal computer passwords (D Mitnick & L Simon 2002).

Social engineering is a major security threat to organizations (Barwick 2012). In order to help organizations successfully manage social engineering threats, it is

crucial for researchers to understand why organizational employees are persuaded to comply with a request posed by an attacker. However, gaining access to individuals' actual behaviour is one constant challenge for researchers in the security field. Little empirical research on social engineering has included real phishing experiments due to ethical concerns related to deceiving participants without debriefing them, and even fewer have been conducted in an organizational setting to understand why organisational employees may or may not fall victim to social engineering. To the knowledge of the authors, only two papers report studies of phishing experiments in an organizational setting (the studies by Jagatic et al. (2007) and Dodgejr et al. (2007) were carried out as real phishing experiments involving university students and not organizational members). One reason for this lack of behavioural studies is that it is challenging to convince organizational managers to participate in studies in which their employees' actual behaviour is being measured. In the experiment conducted by Bakhshi et al. (2009), a phishing mail was sent out to organizational employees as a mean to provide empirical evidence of how many employees succumb to social engineering. The experiment was ceased after approximately 3.5 h. During that period of time, 23 percent of recipients were fooled by the attack. The email included factors related to how the attacker constructs the attack (e.g., trusted e-mail source, attention-grabbing subject, type of social engineering technique used) in order to understand why people fall victim to social engineering. In this paper, we refer to such factors as *external factors.* The results give insight into the problem of social engineering and how vulnerable an organization is to such an attack. No data was, however, collected on personal demographic factors and personal psychological factors to understand why organisational employees succumb to social engineering. We refer to these factors as *internal factors.* The lack of such data makes it difficult to determine personal antecedents of successful social engineering in practice. Furthermore, no information is given on how the management acted during the experiment (did they act according to normal procedures in the event of an attack?). In the study by Workman (2008) a theoretical framework was developed to empirically investigate personal antecedents of successful social engineering. The results revealed that *trust* and *fear* (among others) had significant influence on why people fall victim to social engineering. The data collection was triangulated by collecting data of subjective perceptions of behaviours and conducting objective observations. However, the information given on *external factors* that potentially could affect an individual to succumb to social engineering is limited, and no information is given on how the management acted in the event of a participant reporting his or her suspicion during the experiment. Further, the experiment was conducted over a period of six months, in which both phishing emails and pretext attacks (over the telephone) against each participant were launched two times each week. It is questionable if such an approach reflects an actual attack and using such an attack frequency may both increase the success of the attacks and the awareness of the participants as previous experience of social engineering has shown to improve an individual's resilience against social engineering (Dodgejr et al. 2007). This could potentially bias the results.

In this paper, the general purpose is to extend the understanding of security behaviours in practice by examining reasons why employees fall victim to social engineering. Specifically, we evaluate personal psychological and personal demographic antecedents of successful social engineering and analyse the influence

of adding target information in an attack (important target information could include the name of the targeted organization's CIO in the email). To more fully understand the complexity of security behaviour, empirical data is collected in a multi-method approach by distributing scenario-based surveys under the false pretence of studying "micro efficiency" and conducting unannounced phishing experiments in three organizations. Scenario-based surveys have been used as a technique to assess the security readiness of organizational members in previous studies (Nohlberg 2005). The final specific purpose with this paper is therefore to evaluate if there exists any correlation between how respondents report they would behave in a given scenario in a survey and how they behave in an experiment. The rest of the paper unfolds as follows. In the next section, theory on social engineering is presented. Then, the methodology of the research conducted is presented. The section that follows outlines the results of the empirical tests based on the multi-method approach applied in three organizations. Finally, the findings are discussed and conclusions are drawn.

2. Social engineering

Social engineering consists of techniques used to manipulating people into performing actions or divulge confidential information (D Mitnick & L Simon 2002). Some attackers attempt to persuade individuals with appeals to strong human emotions such as scarcity or excitement, and others focus on adding target information to increase the effectiveness of their attacks. In this study, we want to examine reasons why people succumb to social engineering with an explicit focus on personal psychological and personal demographic factors and the influence of including target information in an attack.

2.1 Individual determinants of successful social engineering

Studies have shown that likeability and trust explain an individual's susceptibility to social engineering. The results of the study by Workman (2008) suggest that an individual that exhibits a greater trust and likeability is easier to deceive. The study also revealed that fear significantly explained why an individual fall victim to social engineering. The author showed that these relationships were significant through a combination of unannounced experiments and questionnaires. *Fear* and *trust* were therefore decided to be included in the present study. Computer self-efficacy or an individual's perception of the ability to perform a computer-related task (Moos & Azevedo 2009; Rhee et al. 2009) has been studied on several occasions and operationalized in various ways, e.g., as general knowledge of computer and as the number of hours an individual spend on a computer each week. This study operationalizes *computer self-efficacy* as the perceived overall knowledge of computers and how many years an individual has used computers in work-related situations. The influence of an individual's risk behaviour has been tested and identified to significantly influence people's susceptibility to social engineering in the study by Sheng et al. (2010). This construct is not specific to IT-related risks, but strives rather to capture general risk behaviour. Social engineering further aims at exploiting human emotions which in turn will affect a person's helpfulness. Therefore, social engineering attacks depend on the natural helpfulness of human users (Luo et al. 2011). In line with these arguments we include *risk behaviour* and

helpfulness in the present study. To capture potential personal demographic antecedents of social engineering, we include *age* and *gender* in the survey instrument. Age and gender have been studied on previous occasions (Workman 2008; Dhamija et al. 2006; Sheng et al. 2010).

2.2 Role of target information

There are a few studies that have estimated the increased effectiveness gained through increasing the amount of target information in a phishing email. In the phishing study carried out by Jagatic et al. (2007), the value of adding target information provided by a social network was estimated. The authors conducted unannounced phishing experiments using students at their university and found that phishing using data provided by social networks gave a 72 % success rate, whereas attacks without such target information gave 16 % success rate. This study is however not very representative to the domain at large as very few phishing attacks are targeted against a specific individual rather than a larger group of individuals. Jakobsson & Ratkiewicz (2006) performed four unannounced experiments on the topic of the online auction system "rOnl". The authors studied the importance of two target specific pieces of data; whether to include the name of the recipient in the email or not ("No name" or "Good name") and the type of link provided in the email ("Good link" or ""evil" IP link" or ""Evil" Subdomain link"). The authors found that the provided link was the most important variable; the presence of recipient name in the email did not have any major influence on the success rate of an attack. This paper, as Jakobsson & Ratkiewicz (2006), aims to study social engineering attacks which are representative to the attack type in general. As a consequence a very important factor is that the attacks need to be automated to some degree. This study utilizes two types of methods to assess susceptibility to social engineering attacks: scenario-based surveys and phishing emails, and these are categorized as one fully automated generic attack (in our case; an attack that only require an email address) and one attack that is targeted against Enterprises in Sweden. The attacks are further described in Section 3. The actual scenarios in the survey and emails used in the experiment can be obtained from the corresponding author.

3. Methodology

As there are characteristic differences between self-perceptions of behaviour and actual behaviour, it is preferable to observe behaviour when possible. However, since all possible behaviours related to social engineering cannot be observed, observation alone is also incomplete (Workman 2008). In order to increase the understanding of complex human behaviour related to social engineering, this study combine these two data collection methods and thus are able to assess how well these two methods go together and if any relations between the results obtained from these two methods exist.

The studies were carried out from April to December 2012. Four different types of tools were used to collect data during this period: scenario-based surveys, experiments, journals and follow-up interviews. In order to conduct phishing experiments, the management of the three organizations was first approached to get

their approval and support for the study. The Chief Executive Officer (CEO) of each organization was notified of the study and the IT manager from each organization took part in designing the study and collecting data. These individuals were also the only employees aware of the study. The nature of the study is such that details about the organizations cannot be presented, but a short description is given as follows: The first organisation has 11 full-time employees (not counting the CEO and IT manager). The company's main focus is in the human resource field in where they conduct employee surveys followed by management coaching and quality improvements. Most of their clients are in the public sector and for the sake of their clients integrity they perceive information security to be important. The second organisation has 32 full-time employees, and the third organization has 49 employees (not counting the CEO and IT manager). Both organisations are in the electrical power domain and the inherent need for security in a these companies is therefore high. All employees of each organization were chosen to be included in the study to maximize the sample size and thus having 92 participants in the study.

3.1 Scenario-based Survey

A smoke-screen approach was used as previous research has argued that it is more effective to capture the employee's security awareness if they are not aware of their awareness being assessed (Nohlberg 2005). This is because the respondents might act differently if they knew that their awareness (or possible lack of) was being assessed. Thus, we wanted to examine if users have a spontaneous awareness of common social engineering cons. The context of the survey was the need to determine how effective an organisation's employees are in the process of performing small work-related tasks during a typically day at work.

The data collection phase started by the IT managers sending out an email and informing their employees about a study in "micro efficiency" and encouraging them to answer a survey related to this study. In line with the purpose of applying a smoke screen approach, three of the scenarios were general scenarios and three where security-related. We attempted to construct scenarios that reflect three actual attacks: update of a well-known software for displaying, printing and managing documents (scenario generic attack), update of the organisation's security software (scenario targeted attack) and acquire of computer password (scenario password). Each scenario was followed by a question to find out what the respondent would do in the outlined scenario on a 7-point Likert scale from 1 to 7 with three fixed points (1: I'll do what I'm asked to do instantly; 4: I hesitate and ask if I can come back to the requester; 7: I completely refuse to do what I'm asked to do). The survey was followed by questions that were aimed to measure the independent variables. To avoid raising any suspicion among the respondents, the questions that were related to measuring the independent variables were explained to be general diagnostic questions and were described to not be related to the scenarios. These dependent variables were measured with single items. Single items are useful and effective for their practical advantages like ease of application and the low costs associated with surveys in which they are used. Further, Bergkvist & Rossiter (2007) found that, in their study, there were no difference in the predictive validity of the multiple-item and single-item measures. The items that were extracted from the study by Workman

(2008) were chosen based on highest loading to their construct and as the constructs were identified to have high composite reliability, the items within the same construct measure the same thing. All items except age, gender and computer experience at work, were measured on a 7-point Likert scale from 1 (Strongly disagree) to 7 (Strongly agree). Gender was used as a dichotomous variable with two states; 1 (Male) and 0 (Female). The items are presented in table 1.

Construct	Item
Trust	Friendly people are usually trustworthy.
Fear	I believe it is important to follow the chain of command.
Risk behaviour	I prefer excitement before a calm and safe every-day.
Computer self-efficacy	I consider myself relatively experienced and skilled with computers.
Helpfulness	I like to help other people.
Computer experience at work	How many years have you worked using a desk-top computer?
Gender	Are you male or female?
Age	How old are you?

Table 1: Survey items

To address common methods bias (CMB) we counterbalanced the order of questions in the questionnaire to discourage participants from figuring out the relationship between the dependent (scenarios) and independent variables that we were trying to establish. Further, the respondent's anonymity reduced the likelihood of bias caused by social desirability or respondent acquiescence (P. M. Podsakoff et al. 2003).

3.2 Experimental design

Two experiments were carried out. Their base scenario was the same. However, the content of the emails were significantly different; the first email being a generic large-scale phishing email, and the second a phishing email with specific information about the target organisation inluded. A pilot study was used to verify that all emails were received by their specified recipients, that the web server was reachable, and that the binary could send data through the firewall. An SMTP server (Postfix) and an HTTP server (Apache) were set up at the research department. The attack was carried out as follows. The SMTP server at the research department sends a "malicious" email to each employee. Every email is outfitted with a unique link to the HTTP server at the research department. An employee clicks on the link in the email and reaches the HTTP server at the research department. The HTTP server was set up to: (i) log user information through a PHP script, and (ii) to automatically serve the "malicious" binary to anyone browsing its contents. The HTTP server sends the "malicious" binary to the employee. This binary did not install anything on the system – it served as a one-time SMTP client. When executed it read the name of the system and the logged-in user, and sent this information to the email account of the conducting researcher (through the mentioned SMTP server at the research department). When the binary had read the system variables and sent these to the researcher it abruptly ended, giving the end-user an error message. The binary also had the correct product icon, but with no specified publisher. The researcher is

notified that the binary has been executed, when it was executed, whom that executed it, and on which system that it was executed.

The attacks in the experiment reflected two scenarios outlined in the distributed survey (generic and targeted attack). The rationale was to enable evaluation of any correlation between how respondents report they would behave in a given scenario and how the behave in the experiment. The first experiment concerned an update of well-known software for displaying, printing and managing documents which was employed on all computers in the enterprise. In this paper, the name of the software is Knylo Reader (the name is obfuscated through ROT10). This product is in the enterprise updated through a service which is installed along with the application. This attack is not targeted at any particular user or organization; from an attackers perspective a recipient is the only information that is required. The domain (www.knldownloads.com) was used to point to the "malicious" HTTP server at the research department. The email was spoofed from support@knylo.com and the user was requested to download the latest version of their software (version 11, which was not released yet at the time of the study). The content of the email was written for the exercise but builds significantly on previous actual phishing attacks using the same product. Furthermore, it was qualitatively reviewed by five external researchers.

The second mail concerns a targeted attack against enterprises in Sweden. The context of the email involves the updating of the enterprise's antivirus software with a temporary add-on as the current antivirus version does not cover the virus that has infected some of the organisation's computer systems. The user is requested to click on a link and update the current antivirus software with the temporary add-on. The email was spoofed from the IT managers' actual email addresses. The whole email was written in grammatically correctly Swedish. This email was specifically written for the experiment and reviewed by five external researchers. However, it was not specifically customized for the studied enterprise, but rather Swedish enterprises in general. In practice there was no need to update the antivirus software at the enterprise. Furthermore, this type of actions is carried out from a central IT-administration, and the IT-managers don't pose any requests of this kind through e-mail. The email was also composed without knowledge regarding how the IT managers typically expressed themselves. As a consequence, the email differed significantly from the style of actual emails by the IT managers. This should serve to make the results of the attack representative to the population at large. The second experiment as such exhibit two critical differences compared to the first experiment: (i) the attacker has to be able to write in Swedish (or consult a third party, e.g., an online service, to translate it) and (ii) the attacker has to find the email to the IT managers of the targeted organisations and be able to relate those individuals to all others in the targeted organisations. In practice, this type of targeted attack is much more effort-demanding than the generic attack. However, such information can be easily accessible; especially for small enterprises. For example, present on the company website or on social media sites such as LinkedIn.

3.3 Analysis Methodology

In order to analyse the relationship between the individual factors and the dependent variable, point-biserial correlation was used. The point-biserial correlation coefficient is a special case of Pearson correlation and can handle dependent variables that are operationlised as scale variables and dichotomous variables. For the dichotomous variable the values typically are 1 (presence) and 0 (absence) (Glass & Hopkins 1995). Thus, this analysis technique fits the purpose of study and we used susceptibility to social engineering (i.e. successful attack) as a dichotomous variable with two states: 1 (Yes) and 0 (No).

4. Results

4.1. Survey results

The survey was sent to the 92 participants of the study. One reminder was sent to non-responding participant after one week. Overall, 54 respondents (59 %) completed the survey. Descriptive results are displayed in table 2. The results indicate that a targeted attack (ScTA) would be most effective and acquiring passwords (ScPW) would be most difficult from an attacker's point of view. Notable are the extremely high mean value of the ScTA and the absolute value of the standard deviation (due to several outliers). A box plot for the descriptive results was analysed and the removal of the outliers yielded a mean value of 7.0. However, we decided to keep the outliers in the descriptive results (presented in table 2) to display the fact that some respondents actually did not report that they would instantly do what they're asked to in a targeted attack.

	Min	Max	Mean	Std. Deviation	N
ScGA	1	7	4.630	2.192	54
ScTA	1	7	6.593	1.125	54
ScPW	1	7	3.926	2.073	54

Table 2: Descriptive survey results

4.2. Experiment

An overview of the results from the phishing experiment can be seen in Table 3. Eight out of 92 recipients, or 8.7 %, clicked the "malicious" link in the generic attack. Three recipients also executed the "malicious" binary. Adding target information in the attack significantly increased the number of employees clicking on the "malicious" link. In the targeted attack, 29 out of 92 recipients, or 31.5 %, clicked the "malicious" link. Six of these individuals executed the "malicious" binary. One individual who executed the binary during the first experiment also executed the binary during the second experiment. Furthermore, all of these individuals executed the binary several times. However, none of the employees executed the binary on more than a single account and system. Nevertheless, these systems could in theory have been more or less vulnerable to the attack during these different executions. During both experiments (the second created a larger amount of activities at the organisations) the IT managers received reports from security-aware

employees. Since the experiment was supposed to be representative to an actual attack in practice and we wanted to capture management behaviour during this event, the IT-managers were told to act as they normally do in an event of a security attack. Therefore the experiment was ceased by the IT-manager sending out a warning about the emails, after approximately 20 minutes in the first attack and after 10 minutes in the second attack. However, there were still employees trying to access the malicious website after the official warning (and knowing that it in fact was malicious). The last attempt to access the malicious website occurred 20 hours after the generic attack and 24 hours after the targeted attack. We can think of two possible reasons that explain this phenomenon: (i) curiosity and (ii) not knowing the dangers involved when browsing malicious websites.

Click link	No.	Percent
Generic attack	8	8.7
Targeted attack	29	31.5
Execute binary		
Generic attack	3	3.3
Targeted attack	6	6.5

Table 3: Overall results from the phishing experiments

4.3. Individual factors explaining susceptibility to social engineering

One of our purposes was to evaluate individual factors that explain why organisational employees succumb to social engineering. Due to the limited sample size associated with the execution of the binary, the analysis was based on individuals clicking on the "malicious" link. Further, we could only use data from the respondents that actually completed the distributed survey (n=54). The results are presented in table 4. ExSA refers to successful attack during the experiment, while ScGA, ScTA and ScPW refer to the three scenarios outlined in the survey described in section 3.1: general attack, targeted attack and password. The statistical results reveal that *computer experience at work*, *gender* (females tend to be less susceptible to a generic attack than men), and *helpfulness* has a significant correlation with behaviour reported by respondents in the scenario-based survey, while *trust* and *risk behaviour* significantly affects the actual behaviour during the phishing experiments.

	ExSA	**ScGA**	**ScTA**	**ScPW**	**N**
Trust	.285*	.092	-.031	-.017	54
Fear	-.070	.096	-.106	.037	54
Risk behaviour	.305*	.079	-.134	-.004	54
Computer self-efficacy	-.010	.210	.008	.018	54
Helpfulness	-.094	.119	-.018	.291*	54
Computer experience at work	.003	-.285*	.204	.087	54
Gender	.043	-.380**	-.108	-.095	54
Age	.144	-.176	.223	.087	54

Table 4: Significance of individual antecedents

Notes: * indicates statistically significant at $p < 0.05$; and ** at $p < 0.01$.

4.4. Combined results

The final purpose was to examine if there exists any correlation between the results from the self-report study (how respondents report they would behave in a given scenario in a survey) and the observations (their actual behaviour in an experiment). We found no correlation between these variables. A correlation matrix of nonsignificant values is presented in table 5 to illustrate the relations between the three social engineering scenarios in the survey (ScGA, ScTA, ScPW) and the actual social engineering attack (ExSA). To examine how the participants reacted to the experiment, we distributed a follow-up survey and when possible, conducted semi-structured follow-up interviews. Overall, the participants perceived the study to be important, had positive feelings about the study and that the study had increased their interest of information security in general and social engineering in particular.

	ExSA	**ScGA**	**ScTA**	**ScPW**
ExSA	1	-.071	.026	-.049
ScGA	-.071	1	-.032	.015
ScTA	.026	-.032	1	.205
ScPW	-.049	.015	.205	1

Table 5: Pearson correlation coefficients for cross correlations

5. Discussions and Conclusions

Social engineering is a major security threat to organizations. One explanation for the threat is the increased effectiveness and robustness of technical security components which has made it more difficult to successfully attack computer systems using purely technical means. A way of compromising information security is then to manipulate computer users into installing malware on their computer or revealing their passwords. In order to understand how to manage social engineering threats, this study tries to understand security behaviours in practice by investigating factors that may cause an individual to fall victim to social engineering. In doing so, this study makes important contributions to the body of knowledge on social engineering in general and reason why organizational employees fall victim to those types of attacks in particular. First, this, to the best of our knowledge, is the first study that has combined a smoke screen survey approach with phishing attacks that are representative to the attack type in general (the attacks are automated to some degree) when collecting data on security behaviours. Second, the results reveal that the degree of target information in an attack increases the likelihood that an organisational employee fall victim to an attack. This results is in line with the results obtained by Jakobsson & Ratkiewicz (2006) and Jagatic et al. (2007). Therefore, we argue that organisations should consider the potential benefits from making enterprise-specific information such as employees' email addresses and titles of organisational members publically available, against the risk that this target information can be collected by an attacker to both spoof e-mail addresses and to instil trust in organisational members. In the end, it's up to the organisation to balance the need to enable the business against the need to secure information assets.

Third, our study has identified that *computer experience at work*, *gender* (females tend to be less susceptible to a generic attack than men), and *helpfulness* showed to have a significant correlation with behaviour reported by respondents in the scenario-based survey, while *trust* and *risk behaviour* significantly affects the actual behaviour during the phishing experiments. These findings indicate that a practical implication could be that organizations should include techniques that are used by social engineers to instil trust and encourage helpfulness and risk behaviour in their security awareness programs.

Our study revealed that surveys and observations capture different factors that explain security behaviours. However, we acknowledge the challenges in collecting data in the critical domain of information security and thus do not rule out one or the other method as we believe they have both advantages and disadvantages. Some might argue that observations capture the actual behaviour. We argue that using observations, exclusively, cannot fully capture the human complex behaviour. For instance, the follow-up survey and interviews revealed that there were occasions in which participants were encouraged and convinced by their colleagues to click on the link. Obviously, this makes it difficult to say if the respondent's actual susceptibility to the attack was measured when appeals to social norms might influence the results. We acknowledge the difficulties in measuring security behaviour and suggest that a deeper understanding of this phenomenon is required, recommending further use of a multiple method approach when attempting to measure security behaviours.

The ethical dilemma related to conducting social engineering experiments in practice makes it rather challenging to recruit participant organisations. Therefore our conclusions are based on relatively few samples. This makes it difficult to generalize the results gained from this study to the domain at large. Nevertheless, it is important to recognize that this study provides insight to properties never before studied. Also, the sample size is comparable to other phishing experiments using unaware respondents (e.g., Jakobsson & Ratkiewicz (2006)). A further limitation is that the scenario-based survey and two experiments were conducted on the same sample of respondents – it is possible that the results of the second experiment are biased from the first. However, there is strong reason to believe that this is not the case: (i) the first experiment (Knylo Reader) was launched 2 months after the survey had been completed and is similar to other spam that is frequently received by employees at the enterprise, and (ii) the second attack was launched approximately three months after the first attack. Finally, we did not spoof a legal website or constructed our own "malicious" website for the study. After the binary had read the system variables and sent these to the researcher it abruptly ended, giving the end-user an error message. We can only speculate the difference in effectiveness if we had spoofed a legitimate website or constructed our own website that serves the purpose of the study.

6. References

Applegate, S.D., 2009. Social Engineering: Hacking the Wetware! *Information Security Journal: A Global Perspective*, 18(1), pp.40–46.

Bakhshi, T., Papadaki, M. & Furnell, S., 2009. Social engineering: assessing vulnerabilities in practice. *Information Management & Computer Security*, 17(1), pp.53–63.

Barwick, H., 2012. Social engineering, big data top security priorities for 2013: Gartner. *Computerworld*. Available at: http://www.computerworld.com.au/article/441539/social_engineering_big_data_top_security_priorities_2013_gartner_/ [Accessed January 10, 2013].

Bergkvist, L. & Rossiter, J.R., 2007. The Predictive Validity of Multiple-Item Versus Single-Item Measures of the Same Constructs. *Journal of Marketing Research*, 44(2), pp.175–184.

D Mitnick, K. & L Simon, W., 2002. *The Art of Deception: Controlling the Human Element of Security*, Indianapolis, Indiana: Wiley Publishing.

Dodgejr, R., Carver, C. & Ferguson, A., 2007. Phishing for user security awareness. *Computers & Security*, 26(1), pp.73–80.

Glass, G. V. & Hopkins, K.D., 1995. *Statistical Methods in Education and Psychology* 3rd ed., Allyn & Bacon.

Jagatic, T.N. et al., 2007. Social phishing. *Communications of the ACM*, 50(10), pp.94–100.

Jakobsson, M. & Ratkiewicz, J., 2006. Designing ethical phishing experiments. In *Proceedings of the 15th international conference on World Wide Web - WWW '06*. New York, New York, USA: ACM Press, p. 513.

Luo, X. et al., 2011. Social Engineering: The Neglected Human Factor for Information Security Management. *Information Resources Management Journal*, 24(3), pp.1–8.

Moos, D.C. & Azevedo, R., 2009. Learning With Computer-Based Learning Environments: A Literature Review of Computer Self-Efficacy. *Review of Educational Research*, 79(2), pp.576–600.

Nohlberg, M., 2005. Social Engineering Audits Using Anonymous Surveys – Conning the Users in Order to Know if They Can Be Conned. In *Proceedings of the 4th Security Conference*. Las Vegas, USA.

Podsakoff, P.M. et al., 2003. Common method biases in behavioral research: a critical review of the literature and recommended remedies. *The Journal of applied psychology*, 88(5), pp.879–903.

Sheng, S. et al., 2010. Who falls for phish? In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*. New York, New York, USA: ACM Press, p. 373.

Workman, M., 2008. Wisecrackers: A theory-grounded investigation of phishing and pretext social engineering threats to information security. *Journal of the American Society for Information Science and Technology*, 59(4), pp.662–674.