To Deceive or Not to Deceive! Ethical Questions in Phishing Research

Rasha Salah El-Din
University of York
York, UK
Rasha@cs.york.ac.uk

Interest in Human factors in phishing has been growing both in HCI and security communities in the past few years. Despite this interest, conducting covert user studies is associated with a number of ethical and legal challenges for phishing researchers. This paper discusses the need for deception, the implications of deceiving and the legal restrictions in terms of phishing study in the UK. We thematically analyzed these implications from the viewpoints of three stakeholders; ethics committees, researchers and professional bodies. Then we provide a roadmap for researchers to get balanced and timely ethical assessment of their proposed research.

Keyword: Ethics, HCI, Research methodologies, Security, Phishing

1. INTRODUCTION

Phishing is a widespread and pernicious practice where criminals seek to obtain money and confidential information such as usernames, passwords or credit card details from people under false pretences. The first step in protecting people from phishing is understanding the dynamics of phishing, the psychology of both the attacker and the victim and analyzing users' decision making strategies reacting to phishing attacks.

Yet, studying the variance in people vulnerability for phishing and reasons behind it is a much under-researched area. In order to study why some people fall for phishing while others do not, we need to understand people's online behaviour. However, obtaining observational data about users' security practices is extremely challenging [4]. When people are aware that their behaviour is monitored, they tend to behave differently than they normally do. And when they are being monitored without being informed, the researchers are accused of breaking a number of laws.

In this paper, we discuss the ethical implications of the use of deception in phishing research. The paper is structured as follows: We first compare the different approaches for studying phishing and the ethical issues they raise. We give particular attention to in-the-wild studies. We start by defining deceptive studies. Then we analyzed the implications of deceptive research from three perspectives; ethics committees, researchers and professional bodies. Finally, we present our view of good practice and outline a roadmap for phishing researchers for designing ethical phishing experiments.

2. BACKGROUND

Generally, there are three main approaches for phishing research; Self-report Studies, controlled lab studies and in-the-wild field studies. The latter is the most ecologically valid, yet the most ethically and legally complicated.

2.1 Self-report Studies

Phishing self-report studies involve the use of questionnaires, online surveys, interviews or polls. Participants are often chosen randomly to answer a set of questions about their past phishing experience, recent losses or latest corruptions of systems and credentials [8].

This research approach has many limitations, one of which is underestimating the risk of phishing if significant number of real phishing attacks were missed and not reported by participants [8]. This happens when victims are either unaware they have been attacked or do not want to reveal they fall for phishing attacks out of embarrassment.
It is also possible that Self-report studies overestimate risk if the participants report non phishing incidents as phishing. This happens as a result of participants' unawareness of what exactly phishing is. An example of that is someone who finds in his credit card bill charges for items he has not purchased. He may suppose this is phishing

and report it as so, while it might be an incident of fraud [8].
Overestimation of phishing risk can also occur if people reported legitimate messages they got from their bank, mobile operator or a real service provider as a phishing attack.

The underestimation or overestimation of phishing risks increase less for interviews more than for polls and online surveys, where there is no direct contact between the researcher and the participants. While in interviews, the researcher can help clarify things for the participants and so gets more precise answers.

Yet, interviews have their own problem when used in security research. People tend to claim to do something, regarding their security practices, but in reality they do something else [4]. One reason is that participants want to impress the researcher and look smarter in front of her. This problem is often referred to as 'the researcher effect'. Here the age, gender or race of the researchers may affect the result they obtain [7].
In this case, these results are ecologically invalid and hence can not be generalized to the real world as they are not a true representative of it.

Another problem associated with self-report studies, is that their non-intrusive nature does not allow the identification of cause and effect [7]. That is why some researchers prefer lab studies to be able to infer causality.

2.2 Controlled-Lab Studies
Lab studies are often used to measure users' ability to detect phishing. It is also called 'Phishing IQ Tests'. It is based on conducting closed lab experiments in which the participants are shown a number of email messages and websites and are asked to distinguish between phishing and legitimate ones.

The main draw back of phishing lab studies is that they are creating an artificial environment that is not similar to that of the real world. It is well known in security research that security practices have rarely been the primary goal of the users, they are not tasks in themselves [11].
Users don not sit down at the computer to "do security" [5], instead, security is an impeded task in other tasks.

Therefore, phishing lab studies are actually imitating a non-natural task that users never perform in real life. Normally, users are not sitting particularly for distinguishing and detecting emails they receive against certain phishing criteria. Alternatively, they deal with phishing while they are performing other activities like checking their emails, navigating through the internet or may be walking in a mall if we are talking about mobile phishing. So isolating users from their daily normal activities to set at a computer just to say which messages they believe are phishing and which are not will result in flawed studies.

Moreover, in a lab study, participants do not feel they are at real risk. They know they are part of a phishing study; both the data and the attack are faked. In an observation made by Whalen and Inkpen about their web security lab experiment [5], the participants did not act to protect the data treat as if it was their own. This means that the knowledge of the existence of the study biases the likely outcome of it [8] the users' real behaviour is not measured.

There is also a possibility that the results of phishing lab studies are affected by 'evaluation apprehension'. This refers to a special type of anxiety that arises when a subject knows he or she is evaluated [2]. That is due to the fact that many participants think the experiments are testing their abilities.

These drawbacks of the self-report and lab studies make it hard to generalize their results to the real world. Looking for a more reliable methodology to help them observe people in more natural settings and at the same time isolating the causal variables studied and ruling out all other explanations of the effect; confound variables, many phishing researchers go for in-the-wild field studies.

2.3 In-the-wild field studies
In this type of studies, researchers simulate a real phishing attack and observe participant's behaviour towards it. In order to do so, researchers need to deceive the participants to the real purpose of the study.

Using deception in research means that researchers deliberately withhold some of the research procedures, mainly its purpose, from the participants aiming to have unbiased conclusions that may result if the participants know they are participating in a phishing experiment.

Not only do these experiments measure the real response to phishing, but they can also measure the threat of attacks that did not occur yet and they can assess the success rates of countermeasures that are not yet deployed.

This approach is the one discussed in this paper.

3. THE ETHICAL CHALLENGES OF DECEPTIVE PHISHING RESEARCH

Acquiring ethical clearance is a mandatory prerequisite for phishing research. Hence in-the-wild field studies need to be reviewed and approved by the ethics committees of the concerned research institutions (RECs). Although this type of research is the most ecologically valid compared to other types explained above, most of these studies are rejected by the ethics committees due to the use of deception [10]. This brings up a debatable question; “Can we deceive users, if our goal is to better understand how they are deceived by attackers?"[3] (P1).

3.1 Ethics Committees’ Perspective

Many ethics committees oppose deceptive withholding of information from research participants. Their concern is based on the grounds that it is unethical, contradicts with informed consent and potentially harmful to the participants [9]; Ethics committees believe it is not justified to deceive people for the interest of research and that generating new knowledge should never override the participants' welfare [1]. Ethics committees also regard deception as a limitation to the participants' control over risk they may be exposed to. Ethics committees are referring here to possible psychological damage or distress.
They also express their fear of possible prosecution by research subjects against the research institution for breaking individuals' rights, invading their privacy and contravening legislation on spam emails and texts.

3.2 Researchers’ Perspective

On the other hand, researchers insist that "given the fact that a piece of research involves deception does not in and of itself make it morally problematic" [1] but rather the rationality behind withholding information from the person being deceived. In other words; the ethics committee's decision should be based on the research context. Precisely as Siber [9] described; "the very strong form of deception can be used in utterly harmless and delightful way" (p.4).
A group of deception proponents presented the participants' point of view. Those advocates introduced what can be referred to as *"The Joy of being deceived"*. They proved that deception has been a source of pleasure to participants. They enjoyed participating in deceptive studies more than non-deceptive ones and were educated more [1]. A willingness to participate again in similar deceptive studies was expressed. Participants also mentioned they found deception unavoidable and some described the experiment's worth as a learning experience and as a scientific endeavour [1].

3.3 Professional Bodies Perspective

Ethics of research are regulated by professional bodies such as The American Psychological Association (APA), the British Psychological Society and Belmont report. Here we provide examples of their views about deceptive research.

In its ethical principles for conducting research with human participants, The BPS literally states that:

> *"It may be impossible to study some psychological processes without withholding information about the true object of the study or deliberately misleading the participants"*

Not only was this the BPS view of deception, but also according to the Belmont report,

1. *Informing subjects of some pertinent aspect of the research is likely to impair the validity of the research.*
2. *In many cases, it is sufficient to indicate to subjects that they are being invited to participate in research of which some features will not be revealed until the research is concluded.*
3. *Care should be taken to distinguish cases in which disclosure would destroy or invalidate the research from cases in which disclosure would simply inconvenience to the investigator.*

As per the APA, it justified the use of deception as a research methodology *as “it enables valid inference by reducing causal ambiguity, or confounding, to a minimum”* [12].

An inspection of these guidelines can ascertain how phishing research lies under the category of research where deception is a necessity.

4. A ROAD MAP TO PHISHING RESEARCH

Our former analysis shows there exists a conflict between the three stakeholders that results in researchers' inability to peruse their deceptive phishing research. Building on this analysis and reflecting from own real world experience of designing mobile phishing experiments; we present a roadmap for phishing researcher to help them acquire balanced and timely ethical assessment for their research.

Each road map dimension defines one important ethical aspect of conducting Phishing deceptive research. The overall roadmap consists of three categories that group together eleven different

dimensions. The phishing roadmap categories and dimensions are as follows with dimensions indented below categories:

- Pre- Launching Phishing simulated attack
  - Preparing fraudulent text
  - Preparing Press release
  - Warning administrative bodies
  - Pre-Informed Consent
  - Statement of Confidentiality
- Launching Phishing simulated attack
  - Data Protection
  - Protecting the Researchers
  - Minding the Participants' Wellbeing
- Post- Launching Phishing simulated attack
  - Debriefing the Participants
  - Post-Informed Consent
  - Data Protection

Each of the eleven dimensions has several indicators that define an ultimate goal for the dimension to ensure ethical conduct of the study. For instance, the fraudulent text should ideally be a simple request and should avoid any element of coercion that could cause anxiety or distress.
Another aspect is warning concerned bodies such as support and administrative staff as participants may report the 'phishing attacks' and this would avoid time spent chasing them up.

In regard to debriefing, researchers need to ensure participants are provided with sufficient information at the earliest stage. Concerning participants' wellbeing, researchers should be aware that the welfare of their participants are minded according to the duty of care law, otherwise, a duty of care is owed to the plaintiff. For that, our road map outlines a procedure that involves consulting appropriately upon the way that the withholding of information or deliberate deception will be received. An 'Anxiety and Panic Handling' training is a must (both Universities' counselling service and health and wellbeing service can assist in that).

On the subject of the legal issues raised by ethics committees and in the absence of an affiliated law that organizes security research in general and phishing research in particular, we give advice to researchers to comply with human rights Act 1998. This includes paying attention to participants' right in privacy and also being aware that the act has to be balanced against the wider public interest and good. In our own mobile phishing experiments, a new 'Pay As You Go' SIM card had to be used and to be dedicated only to the experiment. It had to be kept secured in a locked filling cabinet in a locked room. As soon as the study was finished, the data was deleted prior to the SIM physical destruction.

5. CONCLUSION

Although deception is a well-established research methodology in psychology, it is relatively new to security related research and accordingly provokes ethical debate. We argue that the use of deception in phishing research can be totally safe. We provide a roadmap for researchers to ensure the ethical conduct of phishing experiments. As for the current status of research, we argue the problem is mainly procedural. We call for phishing research proposals to be handled by social science ethics committees rather than physical sciences ones, or at least to invite members from psychology department to give opinion regarding ethical issues raised.

6. REFERENCES

[1] Athanassoulis, N., Wilson, J. When is deception in research ethical. Clinical Ethics 4, 1 (2009), 44-49.
[2] Bagley, P., Evaluation Apprehension: An Examination of Affect in the Audit Environment. 2007.
[3] Dittrich, D., Bailey, M., Towards Community standards for ethical behaviour in Computer.
[4] Dourish, P., Grinter, R., Delgado de la Flor, J., and Joseph, M. 2004. Security in the Wild: User Strategies for Managing Security as an Everyday, Practical Problem. Personal and Ubiquitous Computing, 8(6), 391-401.
[5] Egelman, S., Tsai. J., Cranor, L. Tell Me Lies: A Methodology for Scientifically Rigorous Security User Studies. Workshop on Studying Online Behaviour at CHI'10. April 2010.
[6] Epley, N. & Huff, C. Suspicion, affective response, and educational benefit as a result of deception in psychology research. Personality and Social Psychology Bulletin (1998), 24:759–68.
[7] Field, A., Hole, G., How to design and report experiments. London, UK:SAGE, 2003.
[8] Jakobsson, M., Finn, P. Designing and conducting phishing experiments (2007). IEEE Technology and Society Magazine, Special Issue on Usability and Security.
[9] Sieber, J. E. Kinds of deception and the wrongs they may involve. IRB: A Review of Human Subjects Research, 4, 9 (1982) 1-5.
[10] Soghoian, C. Legal risks for phishing researchers. In eCrime Researchers Summit (2008), 1–11. IEEE.
[11] A. Whitten and J.D. Tygar, Why Johnny Can't Encrypt: A Usability Case Study of PGP. 5.0. Proceedings of the 8th USENIX Security Symposium, August 1999.
[12] Lawson, E. Deception in experimental research: The question of methodological justification. Doctoral dissertation, (1988).University of Sydney, Sydney, NSW, Australia.