Computer Security Research with Human Subjects: Risks, Benefits and Informed Consent

Maritza L. Johnson, Steven M. Bellovin, and Angelos D. Keromytis

Columbia University, Computer Science Department
{maritzaj,smb,angelos}@cs.columbia.edu

Abstract. Computer security research frequently entails studying real computer systems and their users; studying deployed systems is critical to understanding real world problems, so is having would-be users test a proposed solution. In this paper we focus on three key concepts in regard to ethics: risks, benefits, and informed consent. Many researchers are required by law to obtain the approval of an ethics committee for research with human subjects, a process which includes addressing the three concepts focused on in this paper. Computer security researchers who conduct human subjects research should be concerned with these aspects of their methodology regardless of whether they are required to by law, it is our ethical responsibility as professionals in this field. We augment previous discourse on the ethics of computer security research by sparking the discussion of how the nature of security research may complicate determining how to treat human subjects ethically. We conclude by suggesting ways the community can move forward.

Keywords: security research, human subjects, responsible conduct, ethics review committee, institutional review board

1 Introduction

Computer security research frequently entails studying real computer systems and their users. Studying deployed systems is critical to understanding real world problems, so is having would-be users test a potential solution. Oftentimes obtaining these data means interacting with a user, or measuring some aspect of their device. For example, data collection could require installing monitoring software on a user's personal device, instrumenting a website, or conducting a laboratory study. In many cases computer security researchers are doing human subjects research, which is obvious if there is direct interaction with a user, but may also be the case if the collected data was generated by a human. Regardless, it is important for researchers to consider the relationship between the users and the research to ensure the ethical treatment of users.

In this paper we focus on three key concepts in regard to the ethical treatment of users: risks, benefits, and informed consent. These concepts have been used to evaluate the ethics of research in other disciplines and were introduced by the Declaration of Helsinki. They are also widely used by ethics review committees and institutional review boards (IRB). Risk refers to the possibility that something negative will happen to the user as a result of the research. Benefits can be viewed as a something that could positively affect the user, or positively affect a larger population that the user is a member of. Informed consent typically means that the purpose and process of the research are explained to the user, along with the risks and benefits, to allow them to make the decision whether to participate. Researchers may be legally obligated to consider these concepts depending on their location and the nature of the research. In the United States, for example, human subjects research must be evaluated by an IRB, a committee tasked with ensuring people are treated ethically.

As computer security researchers, regardless of whether a committee review is required, we should explore what these concepts mean in regard to our research. Our goal with this paper is to identify some areas for future discussion, argue why our community should take the lead on these concepts, and suggest initial first steps. In the context of computer security research ethics, this paper is concerned with the ethical treatment of human subjects; though the discussion of informed consent, risks, and benefits may be applicable to computer security research that does not involve human subjects in a traditional sense.

Prior work has mentioned this topic as an important piece of the larger discussion of computer security ethics [5]. It's been suggested that perhaps IRBs and ethics committees are in a better position than program committees to provide external ethical review of research [2]. While it may be true that program committees are ill-equipped to conduct an ethical review, based on timing and expertise, turning the issue to the IRB is not an ideal approach. We suggest the community establish best practices for doing human subjects research, similar what has been suggested for vulnerability research [13],

The Ethical Impact Assessment (EIA) framework was introduced to guide the process of determining the potential risks and benefits for stakeholders [10]. The framework is motivated by the same guiding principles that have been used in medical and psychological research [14]. The EIA is a useful starting point for bringing concepts like informed consent and beneficence to the attention of researchers. This paper contributes to the discussion by encouraging researchers to consider how computer security research is the same as medical and psychological research, and how it is different. Exploring these questions will help researchers attain a better understanding of how to apply the concepts of informed consent, risks, and benefits. Usable security researchers have relevant experience, most have interacted with an IRB and have at least a basic understanding of the application of these concepts. They are also able to use their own research as case studies to understand how the research compares to other fields [4]. Since an ethics course is rarely a required part of computer science curriculum, descriptions of how to design a study and how to work with an IRB are instructive [7], as are descriptions of what qualifies as human subjects research [8].

2 Computer Security Research with Human Subjects

As researchers it is in our best interests to determine how risks, benefits, and informed consent apply to our research. We have the deepest knowledge of the area, however, may not have sufficient experience in applied ethics to immediately determine suitable guidelines. We ought to leverage other fields when possible, since this is an issue for other disciplines as well. A step toward achieving this goal is to understand how our research compares to other fields.

To continue the discussion of computer security research with human subjects we ought to compare and contrast our field with medical and behavioral research, the two primary fields of human subjects research. To give a few examples of how our research may differ, in our research there may be the need to collect large amounts of potentially sensitive data, observe login credentials, actively attack the subject, or obfuscate the true purpose of the study [4]. It is reasonable to ask whether our research is different in practice, since many of these examples appear to be quite similar to medical or psychology research. A question that ought to be addressed directly.

Ethics committees and IRBs are tasked with protecting the welfare of human subjects, this includes evaluating whether subjects are sufficiently informed of the risks and benefits of the research, whether the potential risks have been minimized as much as possible, and if expected benefits outweigh the potential risks. Additional factors are considered, but these represent most of the largest concerns. Given that this is an area of expertise for IRB members, but not necessarily for researchers, why would we suggest our community take an active role in discussing how these terms apply to our research? IRBs clearly have expertise in areas that security researchers do not, but it would be a mistake to rely on the existing structure to be the primary source of ethical guidance.

We should look beyond the IRB because, we conjecture, few IRBs have a member with sufficient technical expertise to thoroughly review computer security research. IRBs have deep roots in medical research, other fields that conduct human subjects research have a history of attempting to distinguish themselves from medical research [9, 16]. Many institutions have responded by creating a non-medical IRB. However, given the nascency of security research with human subjects, and the wide array of expertise IRBs are expected to have, it's unclear how many IRBs have adapted their membership to include the necessary expertise.

2.1 Risks

Determining the continuum of risks that may be present in computer security human subjects research is critical, and may benefit ethical decision making for other areas as well. Comprehension of the risks involved is an essential part of IRB review, and is also essential to the primary schools of ethics, consequentialism and deontological. Due to the medical origins of regulations guiding human subjects research, behavioral science researchers have aimed to distinguish themselves from biomedical researchers. Behavioral researchers have asserted that the risks involved in their studies tend to be qualitative, compared to the physical nature of biomedical research [12]. The types of risks include physical, psychological, social, economic, legal, and dignity [15]. Computer security research is more like behavioral research in the sense that the risks typically aren't physical, and can be difficult to quantify and to describe.

In order to set forth a continuum of risks, we need to understand the extremes: what are the characteristics of research that involves minimal risk, and what are the characteristics of research the poses the greatest risks?

2.2 Benefits

Expected benefits of the research ought to be considered in terms of the human who directly participated, as well as the potential benefits to the general users. In medical research, the participants may benefit from participating in a clinical trial for a condition they have, especially when effective treatment is not otherwise available. The direct benefit of psychology research sometimes includes a better understanding of oneself. Computer security studies seem to be more aligned with psychology research, where self-education can be a major benefit of participation. However, the benefit of knowing more about computer security may prove to be quite useful, like when knowledge such as how to avoid a phishing attack can be imparted [11].

2.3 Informed Consent

Informed consent has two primary facets, the first is that the participant is presented with the potential risks and benefits of participation, the second is that the participant is given an opportunity to decide whether to participate. Important differences may exist for our field with the first aspect. Empirical studies have shown that the typical user has an incorrect mental model of basic security primitives [17], and the execution of common attacks like phishing. If they are asked to install monitoring software on their personal device, can they be expected to properly evaluate the risks of participating unless the potential risk is very clearly explained in layman terms? IRB review evaluates whether the consent form is understandable to potential subjects, how do we ensure that both parties comprehend the necessary details? Is a text-based consent form effective? Researchers from other fields have attempted to evaluate the effectiveness of various mediums [1]. In some cases it may be useful to engage in a conversation where the researcher explains key ideas and the participant can ask questions, or to include a brief quiz to gauge comprehension [6].

In some cases disclosing the research purpose in the consent form may threaten the validity of the results. For example, if a researcher plans to study how users respond to an attack, or measure a user's security mindedness, revealing the purpose of the study will influence the participant's behavior. To avoid this researchers can request a waiver of informed consent, or obfuscate the true purpose of the study. Obtaining a waiver typically requires demonstrating that the potential risks are minimal and that other study designs will not suffice. If IRB review is required, the IRB will sometimes request that participants are debriefed once the study is completed, this can serve as a tool to reduce the perceived risks and to ensure the participants questions are answered.

Debriefing takes place after the person has completed the study, it is an opportunity for the researcher and subject to discuss the study and perhaps the true nature of the study. A waiver for debriefing can be granted if revealing the true research protocol may cause the participant distress, and there is minimal risk involved [7]. In our field, debriefing can be an opportune time to increase the benefits of participation by providing the participant with security education. Particularly when participants are being attacked or are answering questions related to their security knowledge and practices. However, it can be difficult to design an effective debriefing message, especially when users participate remotely and are not present in person. Depending on the research topic, the researcher may be in the position to give advice that is known to be effective [11], or they may feel debriefing will raise more concerns than it is able to effectively address thus causing unnecessary distress to participants [7]. It would be useful to have guidelines to help a researcher decide when each technique is appropriate or desirable, perhaps it depends on the amount of risk involved.

3 Moving Forward

This paper raises more issues than it addresses; in this section we will suggest ways that the community can make progress in this area. The first of which is to continue identifying the similarities and differences between our field and fields that have a history of conducting human subjects research. This includes working toward an understanding of the continuums of risks and benefits.

We recommend empirically evaluating our suspicion that most IRBs are unprepared to review research protocols in our field. This conjecture was formed based on our knowledge of IRB membership, the nascency of security research involving human subjects, and the technical nature of some protocols. A better understanding of the expertise and backgrounds of IRB members, and a survey of their level of comfort reviewing various types of protocols would be useful. The study design could be modeled after Buchanan and Hvizdak's evaluation of IRB concerns with research conducted via the Internet [3]. Additional data that could be collected include measuring IRB experience with reviewing computer security research, the number of protocols computer science departments submit each year, and when the first was submitted. It would also be useful to collect data on the sort of questions that arise when reviewing computer security protocols.

Our community could form a community of researchers who have experience with ethics or the IRB process. Researchers could consult with this board during the early stages of the research, and IRBs could also consult with the committee when they need external assistance for the review of a protocol. [1]

[1] IRB membership 45 CFR 46.107 (2009).

Perhaps we need a repository of IRB protocols or study methodologies to encourage the discussion of ethical decision making. This could increase expertise by allowing researchers to gain an understanding of the tradeoffs that were made during the initial stages of the research. Researchers could also describe any IRB concerns that arose, and how they were addressed.

4 Conclusion

We suggest that it is our community's responsibility to explore concepts such as informed consent, risk, and benefits as they pertain to our research. We selected these concepts as the focus of this early discussion because of the important role they play in the IRB review process and because they are the concepts where our research may diverge from other fields. We assign the task to our community because the alternative is to wait for an outside body to impose regulations. The expertise of IRBs and their members will serve as a useful guide, but we must use our intimate knowledge of the domain to ensure the necessary concepts are satisfactorily explored. Much of this paper is dedicated to research with human subjects, however, an understanding of the risks and benefits associated with this research may benefit the larger discussion of computer security ethics.

The recommended directions for moving forward will advance the discussion and lead to a better understanding of the issues at hand. In this paper we introduce a preliminary set of concerns, and suggest possible next steps. We should continue to explore best practices for our field to ensure the ethical design of research methodologies, borrowing from fields where similarities can be found and identifying pertinent differences.

Acknowledgments. The authors would like to thank the anonymous reviewers for their feedback. This work was supported by the NSF through Grant CNS-09-14845. Any opinions, findings, conclusions or recommendations expressed herein are those of the authors, and do not necessarily reflect those of the US Government or the NSF.

References

1. Agre, P., Campbell, F.A., Goldman, B.D., *et. al*: Improving informed consent: The medium is not the message. IRB: Ethics and Human Research 25(5), S11 – S19 (Sep – Oct 2003)
2. Allman, M.: What ought a program committee to do? In: WOWCS'08: Proceedings of the conference on Organizing Workshops, Conferences, and Symposia for Computer Systems. pp. 1–5. USENIX Association, Berkeley, CA, USA (2008)
3. Buchanan, E.A., Hvizdak, E.E.: Online survey tools: Ethical and methodological concerns of human research ethics committees. Journal of Empirical Research on Human Research Ethics: An International Journal 4(2), 37 – 48 (June 2009)
4. Cranor, L.: Ethical concerns in computer security and privacy research involving human subjects. In: FC. pp. 247–249. LNCS, Springer Berlin / Heidelberg (2010)

5. Dittrich, D., Bailey, M., Dietrich, S.: Towards community standards for ethical behavior in computer security research. Tech. Rep. 2009-1, Stevens Institute of Technology (April 2009)
6. Ess, C., AoIR: Ethical decision-making and Internet research: Recommendations from the ethics working committee. `http://aoir.org/reports/ethics.pdf` (2002)
7. Finn, P., Jakobsson, M.: Designing and conducting phishing experiments. In: IEEE Technology and Society Magazine, Special Issue on Usability and Security (2007)
8. Garfinkel, S.L.: IRBs and security research: myths, facts and mission creep. In: UPSEC'08: Proceedings of the 1st Conference on Usability, Psychology, and Security. pp. 1–5. USENIX Association, Berkeley, CA, USA (2008)
9. Gunsalus, C.K., Bruner, E., Burbules, N., Dash Jr., L.D., Finkin, M.W., Goldberg, J., Greenough, W., Miller, G., Pratt, M.G.: The Illinois White Paper - Improving the System for Protecting Human Subjects: Counteracting IRB Mission Creep. Qualitative Inquiry 13(5), 617–649 (2005)
10. Kenneally, E., Bailey, M., Maughan, D.: A framework for understanding and applying ethical principles in network and security research. In: FC. LNCS, vol. 6054, pp. 240–246. Springer, Heidelberg (2010)
11. Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M.A., Pham, T.: School of phish: a real-world evaluation of anti-phishing training. In: SOUPS '09: Proceedings of the 5th Symposium on Usable Privacy and Security. pp. 1–12. ACM, NY, NY, USA (2009)
12. Labott, S.M., Johnson, T.P.: Psychological and social risks of behavioral research. IRB: Ethics and Human Research 26(3), 11–15 (May – June 2004)
13. Matwyshyn, A.M., Cui, A., Keromytis, A.D., Stolfo, S.J.: Ethics in security vulnerability research. IEEE Security and Privacy pp. 67–72 (2010)
14. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research: The Belmont report - ethical principles and guidelines for the protection of human subjects of research
15. National Research Council: Protecting Participants and Facilitating Social and Behavioral Sciences Research. National Academies Press, Washington D.C. (2003)
16. White, R.F.: Institutional review board mission creep: The common rule, social science, and the nanny state. The Independent Review XI(4), 547–564 (2007)
17. Whitten, A., Tygar, J.D.: Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In: 8th USENIX Security Symposium. pp. 169–184 (1999)