

Empirical Research and Research Ethics in Information Security

Edgar Weippl^{1(✉)}, Sebastian Schrittwieser², and Sylvi Rennert¹

¹ SBA Research, Vienna, Austria
`eweippl@sba-research.org`

² JRC Target, FH St. Poelten, Sankt Pölten, Austria
`sebastian.schrittwieser@fhstp.ac.at`

Abstract. Applied and empirical research in information security not only observes and probes technical systems but people are also always involved in these experiments. Therefore ethical considerations are important. Based on our experience and an analysis of well-known papers, we propose the method of ethics case discussions to include ethics considerations in empirical research.

Keywords: Information security · Ethics

1 Empirical Research

Applied research in information security is becoming increasingly important as many large-scale cloud systems and complex decentralized networked systems are used by millions of people today. This complexity is amplified in the engineering of future cyber-physical systems, where the established enterprise perspective of information security needs to be replaced with risk management approaches that balance and trade off needs in safety, reliability, privacy, cybersecurity, and resiliency. In addition, validation of research results is also important for cyber-insurance and aspects of accountability and liability. While many of the applied research conferences happily accept empirical papers, there is too much focus on the reported results and too little emphasis on the methods used to derive the results and build confidence in them.

There are often technical challenges related to the application of empirical research methods to the software development projects. These challenges show the limitations of existing methods when applied to different domain contexts. The use of this research approach in information security is evolving and many of the earlier well-known empirical research findings are hard to reproduce for two main reasons: First, the original data is no longer available (if it ever was) or may have been altered. Second, research ethics have changed and experiments such as, e.g., the social phishing of early days, are no longer an acceptable practice.

We use methods of empirical research and applied computer science:

- We collect data to observe networked systems in order to support the identification of yet unknown and evolving threats.

- Based on these findings, we research methods that can mitigate risks and protect assets in the short and medium term. These methods are evaluated together with our partner companies.
- By identifying the root causes, we work on new approaches that address the weaknesses by fundamentally rethinking existing protocols and implementations.

However, in most cases, people are part of the system that we observe and that we influence through our experiments. For instance, we analyzed how social networks can be attacked [20] or which information is leaked in anonymization networks [19, 21]. Therefore considering ethical implications is important, in particular in cases where users express that privacy is important.

2 Ethics

In the era of “big data”, research based on the enormous amount of data that are more or less easily accessible online—e.g., on social networks, data clouds, and file sharing networks – is becoming increasingly common. The volume of data combined with advances in computing also amplifies the potential ramifications of such research. Therefore, it is important to consider the ethical implications of research that are directly influences real people and real data and to develop ethical principles governing such research efforts.

While in the past, many researchers discussed threats from a theoretical viewpoint (e.g. Thompson et al.’s famous “Trusting Trust” [6] from 1984), the trend is now going towards quantitative analysis of security issues (e.g. [1–5]). From an empirical viewpoint, it would make sense to validate your research into attacks by implementing and testing an attack “in the wild”. From an ethical perspective, however, the situation is more complex and we have to consider two major aspects: Can the results of our research be used to harm others, and do our research activities in themselves harm others?

As for the first question, we have to be aware that our research has applications, some of which we may not approve of. Although it is not always possible to anticipate all the ways in which research can be used, it should be our responsibility to try to think about how the results could be misused, e.g., by oppressive regimes or criminals. Would it be ethical to develop analysis methods for Tor [5] or similar anonymization networks considering regimes might use our work to deanonymize people using them to circumvent censorship?

The second question, which will be the focus of this paper, is one that is familiar from medical and psychological research. While the potential direct impact of security research methods is not usually as dramatic as in those disciplines, the use of unethical methods may nevertheless cause substantial damage. It is all the more important as the line between white hats and black hats can easily become blurred and is only defined by the researcher’s personal ethics.

In light of the growing popularity of empirical computer security research we believe it is important to encourage a serious discussion of ethics in the research community and to develop ethical standards similar to those governing research other disciplines.

2.1 Rules Discussed at Cyber-Security Research Ethics Dialog

The principles discussed in this paper are not directly derived from any particular ethical guideline or borrowed from other scientific disciplines. We base our discussion on the paper of [18]. Our aim is to suggest fundamental ethical principles based on common sense, as we believe it is of the utmost importance to achieve a broad consensus on such principles across the information security community. Without such a fundamental agreement, the development of detailed guidelines or frameworks would be ineffectual. We discuss this idea further in the following section.

Do Not Harm Humans Actively. Today, this is one of the fundamental principles of research and, from a common sense perspective, it seems obvious. However, even disciplines where today it is an immutable principle, such as medical research, have historically ignored this apparently obvious moral imperative (for an example, see the Tuskegee syphilis experiment¹). It was only after decades of unethical medical research where patients were not informed or even actively misinformed about available treatments, not treated at all, or infected on purpose and used as vectors for further infections, that rules were drawn up to define which lines should not be crossed in medical research, such as the Declaration of Helsinki [9]) – a set of ethical principles governing medical research involving human subjects as well as identifiable human material and data², which is widely regarded the cornerstone of medical research ethics.

Today, few would argue against the need for such rules in the field of medical research, where human lives are directly and often dramatically affected. The impact in information security research may not usually be as drastic, but over the past years, there have been studies and experiments that still had a dramatic impact on the individuals involved (often without their knowledge or consent). Although not academic research, the “Craigslist Experiment”² showed what a serious impact unethical studies can have on individuals, and depending on the setup, an academic study on privacy-impacting behavior (perhaps similar to [10]) or cyber-bullying in social networks might have a similar impact on those involved. Ensuring that no humans come to harm in any way must therefore be a priority in academic research in information security. This can be complicated by unpredictable effects on complex systems — how can we ensure that our analysis of, e.g., a botnet does not interfere with the system and its involuntary participants in a harmful way?

Do Not Watch Bad Things Happening. This, again, seems like an obvious principle, but even more so than the previous one, following it can seem like a hindrance to information security research. The underground economy with its spam networks, malware, credit card and identity theft, and a host of other illegal activities is an environment that can generally only be observed in the wild.

¹ <http://en.wikipedia.org/wiki/Tuskegeesyphilisexperiment>.

² The trolls among us. New York Times. Aug 3, 2008. <http://www.nytimes.com/2008/08/03/magazine/03trolls-t.html>.

However, if we observe a botnet without informing the victims or trying to get it shut down, we are standing by as others – network operators, mail service providers, innocent victims who lose money or buy counterfeit pharmaceutical products that may constitute a serious health threat — are harmed.

How, then, can such environments be studied in an ethical way? And would not collecting or deleting identifying data of victims of a botnet absolve one of the moral duty to inform the victims or enforcement agencies? And thinking further, how can we ensure that informing a user that their computer is infected would not cause further harm, e.g. in the form of consequences at work for someone who allowed the workstation to become infected through negligence?

Do Not Perform Illegal Activities to Harm Illegal Activities. A seemingly more complex question is whether it is unethical to use unethical or illegal methods to harm those engaging in illegal activities. For example, a researcher seeking to avoid violating the two previously stated principles, might want to test the effectiveness of botnets by sending spam to prepared test accounts. However, paying for botnet resources would mean funding illegal activity (i.e. maintaining and renting out botnets) with research funds, which is clearly not ethical. To solve this problem, the researcher might decide to use a stolen credit card number and then have it locked, so that the credit card company would revoke the payment. However, even using a stolen credit card would be illegal and unethical. The same goes for botnets: hacking into a distributed network is illegal, whether it is a legal network like SETI@home [12] or Folding@home [13] or an illegal botnet. Unless such an activity is sanctioned by law enforcement, it is not only unethical but also in violation of the law. From a philosophical perspective, this issue may seem complex, but from a legal point of view, it is clear that the end does not justify the means.

Do Not Conduct Undercover Research. This is closely linked to the previous point. Undercover work may require simulation of or actual participation in illegal activities, which is why undercover work by law enforcement officers is strictly regulated (e.g., when it comes to induction rituals in gangs or performing illegal activities as part of an undercover operation). Therefore, some investigations by third parties can only be carried out legally in cooperation with law enforcement. However, in many countries, this is not yet common practice in academic research. Nevertheless, it would be not only unethical but illegal to buy illegal drugs, botnets, stolen credit card numbers or any other illegal goods in the name of science. In addition to the legal ramifications of such activities, researchers should also consider the ethical implications of paying for such goods and thus funneling money into the underground economy.

Despite the general and seemingly common-sense nature of these four principles proposed here, they may seem like serious obstacles to research in the information security domain. If we follow these rules, some empirical studies that previously seemed perfectly acceptable may, in retrospect, fall into a morally grey area or even be completely unethical. This is not unlike the rules governing research in the medical field: while it might be interesting to observe infection

rates and patterns of diseases under real conditions, or compare different therapies the complete absence of treatment, such research is unthinkable today, as it would cause avoidable harm and suffering to humans. Despite these strict ethical rules in medicine and psychology, valuable research is still being conducted. In our opinion, this is what we need in the field of information security: that research that harms or does not help humans becomes unthinkable and we develop new ways of studying important security issues. As these principles can be difficult to implement and it may be daunting to make a decision on the ethics of a given research proposal unaided, we have applied a method from healthcare to security research: ethical case deliberations. The following section describes the background as well as our adaptation of the method for our work at SBA Research.

3 Ethical Case Deliberation

The discussion of ethical issues in security research is relatively recent, particularly in the German-speaking region. Although some articles have been published on the topic in the Anglophone countries and some efforts have been made by security researchers to develop guidelines for research [4, 24], we are still a far cry from the situation in the field of medicine, where ethics committees, ethics guidelines, and ethical case deliberations are par for the course [22, 23].

This paper focuses on ethical case deliberations and how this tool could be adapted to security research. We believe this method could benefit security researchers, who often work it with sensitive personal data, such as private messages, photos, and documents. The data owners have the right to expect an ethically sound treatment of their data, whatever the setting.

If we apply the method of ethical case deliberation, researchers would present their project to an audience of colleagues, lawyers, non-experts, etc., and discuss its ethical implications with them. This allows them to gather opinions, questions, and reservations from people with different perspectives at backgrounds in order to make an ethical decision. Doubts regarding the way in which research is conducted and legal questions should be raised in these deliberations.

To our knowledge, the method of ethical case deliberations has so far never been used in security research, making this a completely novel approach. In this section, we will first lay out the background of ethical case deliberations, then show how we plan to adapt them to the field of security research, and finally, describe an example of such a deliberation and its results at our research Institute.

3.1 Background

Ethical case deliberations come from the healthcare sector, where they support clinical staff in making ethical decisions in their work. Should ethical questions or uncertainties arise in the course of the treatment of a patient, or should there be conflicting opinions about the welfare of a patient and the – e.g., between the

nursing staff and physicians – an ethical case deliberation is convened. It can help reach decisions on whether or not to administer or continue a treatment and other ethically controversial situations [26]. The participants of an ethical case deliberation are people interested in joint decision-making [26, 27] – mainly the clinical staff involved, but depending on the case, it may also include the patient or their relatives. The discussion is moderated by a neutral third party [26]. One frequently used method is the Nijmegen method, but there are also others that can be used [26, 28].

The Nijmegen method consists of four consecutive steps: In the first two, the ethical dilemma as presented and the facts about all relevant aspects are collected. This is followed by an evaluation of the ethical dilemma taking into account all the facts based on ethics guidelines, and finally, the participants reach a joint decision [28, 29].

It has proven useful to base the evaluation and decision on set of ethics guidelines from the medical field, such as the “Principles of biomedical ethics”, which consist of the four principles of respect of autonomy, non-maleficence, beneficence, and justice [22].

The outcome of an ethical case deliberation should be seen as a good decision in as specific case under certain conditions, but it should not be generalized: a change in circumstances, such as other results from other research or new pharmaceuticals could change that situation drastically. In a clinical setting, the results of an ethical case deliberation should be considered a good possible decision that is supported by those involved, but the final decision still rests with the attending physician.

3.2 Approach

Our objective is to adapt the method of ethical case deliberations to the field of security research. This requires a number of changes and decisions regarding participants, moderation, ethics guidelines, method, time frame, and logging. Our recommendations are provided as a starting point only – individual security researchers and the community as a whole can need to come to find a method that works for them. Our recommendations have been tested and appear promising, but alternatives are certainly possible and, where they appear sensible to us, we have noted them.

Participants. We would suggest having no more than ten participants: the more people are involved in the discussion, the longer it gets and the harder it becomes to keep track of individual points. We found six to be an ideal number, as this allows a diversity of opinions without becoming too confusing.

The list of participants should, of course, include the researcher(s) whose project is to be evaluated, as well as their peers. Other important participants are representatives of an ethics committee on security research, if there is one, and a lawyer, should there be any question about the legality of methods used in research. If the proposal is not too technically complex, it could be interesting to also hear the views of laypeople who might be affected by the research.

A balanced ethical case deliberation group should consist of a good mix of internal and external, female and male, as well as expert and non-expert participants. However, it should be noted that recruiting laypeople may delay the start of the discussion, and may therefore not always be an option.

Moderation. The discussion should be moderated by a neutral third party. This person should have experience in the moderation of discussions, particularly when it comes to communication skills and comp tenses in process control, and they should be able to move the discussion forward. The moderator should have at least a basic knowledge of the topic at hand in order to be able to follow the discussion and document the opinions and results (e.g., on a flipchart or whiteboard).

Ethics Guidelines. It would make no sense to use ethics guidelines from the medical field for ethical case deliberations in security research, as the requirements and risks are very different, as discussed in the previous section. However, the four principles we suggested there are very general, making a detailed discussion of a case hard. Therefore, we decided to use the principles defined in the Menlo Report, which were developed specifically for information and community technology (ICT) research [24]. The four core Menlo Principles are respect for persons, beneficence, justice and respect for law and public interest. Each of these principles is divided into several subsections where all aspects of the principles and how to apply them are described in detail. Although there are other suggestions and drafts for ethics guidelines in ICT research, none of them have the scope or specificity of the Menlo Report [30]. Additionally, it is a reasonable assumption that the authors of the Menlo Report will continue to develop and update it.

Method. We suggest the following structure for the discussion: The moderator explains the purpose and method of ethical case deliberations and specifies the time frame for the discussion (see below). If necessary, the participants should introduce themselves (e.g., if an external expert has been recruited who does not know the researchers of the organization). If necessary, the moderator can explain the rules of the discussion, e.g., that only one person should speak at a time or that participants should address each other directly. Whether or not any rules are introduced, it is important to stress that every contribution is useful.

After the moderator explains the process, the actual discussion can start. We found it useful to first present the ethical dilemma and then evaluate it using the four Menlo Principles as a guideline, discussing the dilemma with respect to one principle after the other. If there are severe disagreements concerning one principle, the research question can be changed and the process repeated. After the discussion of the ethical aspects is completed, the participants should try to find a joint decision. All arguments presented should be included in the documentation which, like the decision, should have the support of all participants. At the end of the meeting, moderator should thank everyone for their participation.

We have designed a short presentation guiding the ethical case discussion in SBA Research, which you may adapt freely for your purposes.

Time Frame. We recommend a discussion time of three hours, as this allows plenty of time to discuss all aspects in detail and come to an agreement. Should a conclusion be reached before the end of the allotted time, the discussion can naturally end early.

Logging. It is important to document the process and results of the ethical case deliberation. Depending on the purpose of the discussion, this could be done either by a person not involved in the discussion taking minutes or by recording the entire discussion and using the audio file as a log of the results.

References

1. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G.M., Paxson, V., Savage, S.: Spamalytics: an empirical analysis of spam marketing conversion. *Commun. ACM* **52**(9), 99–107 (2009)
2. Stone-Gross, B., Cova, M., Cavallaro, L., Gilbert, B., Szydlowski, M., Kemmerer, R., Kruegel, C., Vigna, G.: Your botnet is my botnet: analysis of a botnet takeover. In: *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pp. 635–647. ACM (2009)
3. Jagatic, T.N., Johnson, N.A., Jakobsson, M., Menczer, F.: Social phishing. *Commun. ACM* **50**(10), 94–100 (2007)
4. Bilge, L., Strufe, T., Balzarotti, D., Kirda, E.: All your contacts are belong to us: automated identity theft attacks on social networks. In: *Proceedings of the 18th International Conference on World Wide Web*, pp. 551–560. ACM (2009)
5. McCoy, D., Bauer, K., Grunwald, D., Kohno, T., Sicker, D.: Shining light in dark places: understanding the tor network. In: Borisov, N., Goldberg, I. (eds.) *PETS 2008*. LNCS, vol. 5134, pp. 63–76. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70630-4_5](https://doi.org/10.1007/978-3-540-70630-4_5)
6. Thompson, K.: Reflections on trusting trust. *Commun. ACM* **27**(8), 761–763 (1984)
7. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., Savage, S.: Spamalytics: an empirical analysis of spam marketing conversion. In: *Proceedings of the 15th ACM Conference on Computer and Communications Security*, pp. 3–14. ACM (2008)
8. McCoy, D., Pitsillidis, A., Jordan, G., Weaver, N., Kreibich, C., Krebs, B., Voelker, G., Savage, S., Levchenko, K.: Pharmaleaks: understanding the business of online pharmaceutical affiliate programs. In: *Proceedings of the 21st USENIX Conference on Security Symposium*, p. 1. USENIX Association (2012)
9. Kimmelman, J., Weijer, C., Meslin, E.: Helsinkidiscords: FDA, ethics, and international drug trials. *Lancet* **373**(9657), 13–14 (2009)
10. Plc, S.: (2007) Sophos facebook ID probe shows 41 percent of users happy to reveal all to potential identity thieves. <http://www.sophos.com/pressoffice/news/articles/2007/08/facebook.html>. Accessed 07 Feb 2013
11. Commtouch: Internet threats trend report (2012)
12. Anderson, D.P., Cobb, J., Korpela, E., Lebofsky, M., Werthimer, D.: Seti@home: an experiment in public-resource computing. *Commun. ACM* **45**(11), 56–61 (2002)

13. Beberg, A.L., Ensign, D.L., Jayachandran, G., Khaliq, S., Pande, V.S.: Folding@home: lessons from eight years of volunteer distributed computing. In: *Parallel & Distributed Processing, IPDPS 2009* (2009)
14. Wondracek, G., Holz, T., Platzer, C., Kirda, E., Kruegel, C.: Is the internet for porn? An insight into the online adult industry. In: *Proceedings (online) of the 9th Workshop on Economics of Information Security*, Cambridge, MA (2010)
15. Thomson, J.A., Itskovitz-Eldor, J., Shapiro, S.S., Waknitz, M.A., Swiergiel, J.J., Marshall, V.S., Jones, J.M.: Embryonic stem cell lines derived from human blastocysts. *Science* **282**(5391), 1145–1147 (1998)
16. Dittrich, D., Kenneally, E.: *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*, US Department of Homeland Security (2011)
17. Bailey, M., Dittrich, D., Kenneally, E., Maughan, D.: The Menlo report. *IEEE Secur. Priv.* **10**(2), 71–75 (2012)
18. Schrittwieser, S., Weippl, E.: Ethics in security research – which lines should not be crossed? In: *Proceedings of the Cyber-Security Research Ethics Dialog & Strategy Workshop (CREDS 2013) at IEEE Symposium on Security and Privacy (S&P)*. IEEE, San Francisco, May 2013
19. Winter, P., Köwer, R., Mulazzani, M., Huber, M., Schrittwieser, S., Lindskog, S., Weippl, E.: Spoiled onions: exposing malicious tor exit relays. In: Cristofaro, E., Murdoch, S.J. (eds.) *PETS 2014*. LNCS, vol. 8555, pp. 304–331. Springer, Cham (2014). doi:[10.1007/978-3-319-08506-7_16](https://doi.org/10.1007/978-3-319-08506-7_16)
20. Huber, M., Mulazzani, M., Weippl, E.: Who on earth is “Mr. Cypher”: automated friend injection attacks on social networking sites. In: Rannenberg, K., Varadharajan, V., Weber, C. (eds.) *SEC 2010*. IAICT, vol. 330, pp. 80–89. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15257-3_8](https://doi.org/10.1007/978-3-642-15257-3_8)
21. Huber, M., Mulazzani, M., Weippl, E.: Tor HTTP usage and information leakage. In: De Decker, B., Schaumuller-Bichl, I. (eds.) *CMS 2010*. LNCS, vol. 6109, pp. 245–255. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-13241-4_22](https://doi.org/10.1007/978-3-642-13241-4_22)
22. Beauchamp, T.L., Childress, J.F.: *Principles of Biomedical Ethics*. Oxford University Press, Oxford (2001)
23. Dörries, A.: Mixed feelings: Physicians’ concerns about clinical ethics committees in Germany, *HEC Forum*, no. 15, pp. 245–257 (2003)
24. Dittrich, D., Kenneally, E., Bailey, M.: *Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report*, US Department of Homeland Security (2011)
25. Carpenter, J., Dittrich, D.: Bridging the distance: removing the technology buffer and seeking consistent ethical analysis in computer security research. In: *Digital Ethics: Research and Practice* (2012)
26. Dörries, A.: The 4-Step approach - ethics case discussion in hospitals. *Diametros* **22**, 39–46 (2009)
27. Hahn, B., Schulz, M., Hansen, U., Stöcker, R., Kobert, K.: Ethische Fallbesprechungen als Instrument in Psychiatrischen Kliniken - Erfahrungen aus der Klinik. In: *Depressivität und Suizidalität – Prävention, Früherkennung, Plegeinterventionen, Selbsthilfe*, Unterostendorf, Ibicura, pp. 76–83 (2010)
28. Steinkamp, N., Gordijn, B.: Ethical case deliberation on the ward. A comparison of four methods. *Med. Health Care Philos.* **6**, 235–246 (2003)
29. Steinkamp, N., Gordijn, B.: Die Nimwegener Methode für ethische Fallbesprechungen. *Rheinisches Ärzteblatt* **5**, 22–23 (2000)
30. Carle, S.: *Crossing the Line — Ethics for the Security Professional*. SANS Institute (2003)

Information Systems Security and Privacy
Second International Conference, ICISSP 2016, Rome,
Italy, February 19-21, 2016, Revised Selected Papers
Camp, O.; Furnell, S.; Mori, P. (Eds.)
2017, XII, 215 p. 100 illus., Softcover
ISBN: 978-3-319-54432-8