Teaching Information Security students to "Think thief"
Pieter Hartel and Marianne Junger, University of Twente
(Version 4, 15 July 2012)

Key words: K.3 COMPUTERS AND EDUCATION; K.4.1 Public Policy Issues (Abuse and crime involving computers); K.6.5 Security and Protection;

Abstract
We report on an educational experiment where information security master students were encouraged to think out of the box. Instead of taking the usual point of view of the security engineer we challenged the students to take the point of view of the motivated offender. We report on the exciting ideas our students came up with, and on the lessons we learned in designing the course.

Introduction
Some products and services have criminogenic properties that might have been avoided if the designers had also taken point of view of a motivated offender. For example, a beer glass can be dangerous a weapon but a beer glass made from laminated glass shatters like a car window, making it hard to hurt someone [McG10a]. Taking the point of view of a motivated offender is called "thinking thief" [Gam03]. We have studied a number of Information Security curricula and found that some courses devote time to "thinking thief", for example cryptanalysis (mathematics), kernel hacking (systems), and red-blue teaming (networks). A limitation of these courses is that they all focus on technical attacks. However, information security problems in the real world are increasingly caused not by technical attacks but by social engineering, for example the Spear-phishing attack on RSA Inc. in the spring of 2011 [Amo11]. We therefore created an information security course entitled "Cyber Crime Science", which focuses on the social aspects of thinking thief. The purpose of this short paper is to present our experience in teaching the course.

Related work
A penetration testing course would normally include social aspects too, but pen-testing courses tend to lack academic credibility because the underlying theory is missing. In a forthcoming article, Conti and Caroland [Con12] show a variety of ways of teaching students how to think creatively about adversaries. We share the same ideas, but our course is quite different. For example our course is based on the theories of crime science [Jun12], an interdisciplinary field of the social and the physical sciences. Crime science studies everything that is relevant to a crime event in a scientific manner, offering the theories and research methods necessary to practice thinking thief.

Challenge
What we wanted to do is to develop a course that asks information security students to take the point of view of the offender, and to do so in a scientific manner. Practically this meant designing and executing a crime science experiment, where the students had to study the literature on their chosen topic, write a research proposal, build the necessary web sites, tools, and services, recruit subjects, collect and analyse the data, and finally they had to write and present a paper. A quality crime science experiment takes considerable time and skills to prepare and execute. The time requirements range from months to years and the skill requirements necessitate a bachelor degree the social sciences. We were therefore faced with two problems. Firstly, our information security course is a typical one

semester course (6 European credits of 30 hours each). Even in teams of 2 or 3 students, the time available is not sufficient for a rigorous experiment that allows for a publication quality analysis. Secondly, our information security students do not have a formal training in social science research methods. Hence we were faced with the challenge to develop a "lightweight" approach to crime science for information security master students.

Approach

Thirteen teams of 2 or 3 information security master students attended the course. We suggested a range of papers from the literature as a source of inspiration on topics such as war driving [Ber04], anti-phishing training [Kum09], and botnet infiltration [Kan08]. However, we found that the students largely preferred to develop their own ideas.

In the first few weeks of the course, while we lectured on crime science, the student teams drafted a research proposal. After two rounds of feed-back on the research proposals, the teams were given permission to execute their projects, ultimately resulting in 13 six-page papers that were presented at a half day conference at the end of the course. The time constraints put the students under a lot of pressure, with which they coped admirably. Only three students out of the 37 dropped out and the remaining 34 completed the course, either by doing a survey or an experiment.

A survey is person oriented, and focuses on the differences between subjects. An experiment is situation oriented and focuses on the effect of an intervention. Therefore, the opportunity to "think thief" in the case of a survey is limited to questions such as what kind of people do not secure their Wi-Fi base-stations, or to what kind of people accept Face book friend requests willy-nilly. Experiments push thinking thief further than surveys because experiments assess behavioural change, e.g. to what extent do anti-phishing warnings really work?

Planning and executing an experiment is harder than a survey for a variety of reasons. For example an experiment often involves deception, in the sense that the real purpose of the experiment can only be disclosed at the debriefing stage. This requires the approval of the IRB, subjects to sign an informed consent, and debriefing of the subjects. A survey usually avoids some of these steps.

While we were aware of these difficulties, we felt that the learning experience of executing an experiment would be so much greater than the experience of a survey, that we encouraged our students to design experiments, and accepted surveys somewhat reluctantly. In the end 6 teams performed a survey and 7 teams performed an experiment.

Ethical issues

The student projects took place under the responsibility of the lecturers (c.f. sections 1.8 and 1.17 of the Collective Labour Agreement Dutch Universities, CAO NU), who monitored the experiments closely, ensuring that subjects were treated ethically.

Results

Table 1 summarises the 6 surveys on the left and the 7 experiments on the right. N indicates the number of (human) subjects taking part in the study and C indicates the number of subjects in the control group. In all experiments the number of subjects was too low to observe statistically significant results. In some experiments the students were

unable to include a control group, even though they had planned to do so in their research designs. In all studies the subjects were recruited via a convenience sample, rather than a random sampling of an appropriate population.

Surveys	**N**	**Experiments**	**N**	**C**
1. War driving	0	7. QR-code anti-phishing training	57	12
2. Simulated identity theft	100	8. Illegal download warnings	59	0
3. Routine activity theory in online gaming	222	9. Anti-phishing training	66	35
4. Drive-by-downloads	0	10. Privacy training	67	0
5. Geo tagging	22	11. Interactive trash cans	24	14
6. Tor exit traffic	0	12. Fake friends on Face book	28	0
		13. Lost USB sticks	19	0

Table 1 Summary of the 13 student papers

The appendix provides a short paragraph on each of the 13 student papers but here we should like to mention three ideas that have not received much attention in the literature yet.

The first paper (#7) is based on the observation that email-based phishing is probably becoming less effective after years of anti-phishing campaigns. Thinking thief inspired our students to consider alternatives for the ubiquitous phishing email, i.e. QR codes. These are appearing more and more on posters, advertisements, in magazines and people tend to scan them without paying much attention to the actual site that the QR code leads to. Thinking thief about QR codes suggests that, if misused cleverly, QR codes could become an efficient modus operandi for the next generation of phishers. For example a sticker glued on top of an existing QR code would probably not be noticed by most people, and it would probably have a 100% click through rate. The idea has been proposed before [Kie10] but not researched as far as we know. We are currently working on measures to reduce the dangers of QR code based phishing.

The second paper (#8) concerns a method of studying the risk appetite of Internet users. The idea is to build a web browser that is immune to drive-by-downloads, so that the surfing user can freely visit even the most dangerous places on the Internet. An equivalent experiment in real life would be to send subjects with a team of invisible body guards to the most dangerous places on earth. Our students ran out of time building the tool and could not try it out in an experiment. We hope to be able to use the immune browser for the next edition of our course. There is already some literature on this topic [Veg09].

The third paper (#2) is perhaps the simplest and the most ingenious. What happened is that one of the students had lost his wallet. So he went to the police to report the loss. The police gave him a temporary ID, which he took to his bank to ask for a new card and some money. Thinking thief, the student and his team then investigated how easy it would be to

collect all the information necessary to go through this process for a randomly selected member of staff from the university. It should not be a great surprise that they found it easy to collect most of the information necessary. This paper is a survey in the sense that it collects information on subjects, but it has some features of an experiment because the researchers really did go to the police to obtain a temporary ID.

Conclusion

Our students found "thinking thief" challenging but exciting. We have never seen a cohort of students working so hard on a course, which we believe is due to the fact that the students were free to choose their own topics. We are confident that the students have learned and practiced a valuable skill, but more importantly, we hope that they will never forget the importance of trying to put oneself into the shoes of the offender.

Setting aside the issue of the missing control groups, all 13 student projects were in principle based on a sound methodological design, but due to the time pressure, the results were either not statistically significant or not analysed with sufficient depth. This then gives us the answer to the challenge we set ourselves of developing a "light weight" approach to crime science: simply drop the requirement that the results are statistically significant.

Our course gives equal importance to the social and the technical aspects of information security, which in a sense contextualises the social aspects for technical students. We believe that this is an effective way of teaching technical students social science research methods. This knowledge does not apply to just information security but to a much broader range of technical subjects, such as usability engineering, and human media interaction.

We should like to end with a little speculation. Let us assume that our students are generally smarter than the average offender. Then an equal ability of thinking thief on the part of the offenders and our students as future designers of new products and services should give them a significant advantage in the fight against crime and disorder.

References

[Amo11] D. Amorosi. Data breach spring. Infosecurity, 8(3):6-9, May 2011. http://dx.doi.org/10.1016/S1754-4548(11)70032-8.

[Ber04] H. Berghel. Wireless infidelity I: war driving. Commun. ACM, 47(9):21-26, Sep 2004. http://doi.acm.org/10.1145/1015864.1015879.

[Con12] G. Conti and J. Caroland. Embracing the Kobayashi Maru: Why you should teach your students to cheat. IEEE Security & Privacy, to appear, 2012.

[Gam03] L. Gamman and B. Hughes. Thinking thief - designing out misuse, abuse and criminal aesthetics. The Ingenia Magazine, 15, Feb 2003. http://www.ingenia.org.uk/ingenia/issues/issue15/Gamman.pdf.

[Hol09] T. J. Holt and A. M. Bossler. Examining the applicability of lifestyle-routine activities theory for cybercrime victimization. Deviant Behavior, 30(1):1-25, Jan 2009. http://dx.doi.org/10.1080/01639620701876577.

[Jun12] M. Junger, G. Laycock, P. H. Hartel, and J. Ratcliffe. Crime science: editorial statement. Crime Science, 1:1:1-1:3, Jun 2012. http://dx.doi.org/10.1186/2193-7680-1-1.

[Kan08] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamalytics: an empirical analysis of spam marketing conversion. In 15th ACM Conf. on Computer and communications security (CCS), pages 3-14, Alexandria, Virginia, Oct 2008. ACM. http://dx.doi.org/10.1145/1455770.1455774.

[Kie10] P. Kieseberg, M. Leithner, M. Mulazzani, L. Munroe, L. Schrittwieser, M. Sinha, and E. Weippl. QR code security. In 8th Int. Conf. on Advances in Mobile Computing and Multimedia (MoMM), pages 430-435, Paris, France, 2010. ACM, New York. http://doi.acm.org/10.1145/1971519.1971593.

[Kum09] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. Blair, and T. Pham. School of phish: a real-word evaluation of anti-phishing training. In 5th Symp. on Usable Privacy and Security (SOUPS), Article 3, Mountain View, California, Jul 2009. ACM. http://dx.doi.org/10.1145/1572532.1572536.

[McG10a] C. McGinley and C. Till. Design Out Crime: Using design to reduce injuries from alcohol related violence in pubs and clubs. Design Council, Mar 2010. http://www.designcouncil.org.uk/Documents/Documents/Publications/Crime/DesignOutCrimeAlcohol_Insights_Design_Council.pdf.

[Veg09] H. Vegge, F. M. Halvorsen, R. W. Nergard, M. G. Jaatun, and J. Jensen. Where only fools dare to tread: An empirical study on the prevalence of Zero-Day malware. In 4th Int. Conf. on Internet Monitoring and Protection (ICIMP), pages 66-71. IEEE, May 2009. http://dx.doi.org/10.1109/ICIMP.2009.19.

Appendix

The main idea and the results of the thirteen papers can be summarised as follows:

1. "Open WiFi network availability, an analysis" presents a war driving experiment covering 4277 WiFi access points in three areas of a small city. Over 10% of those were found to offer no security, and only two access points were configured such that on a subsequent scan by Google Street view, they would not be monitored. The researchers had planned to interview a number of access point owners but due to lack of time they did not achieve this (i.e. N=0). In the end this was a purely technical project, and not a real crime science experiment.
2. "How to ruin someone's life in three easy steps" presents a study where the researchers selected the names of 100 potential targets randomly from the University telephone directory, then collected the data necessary from the Internet to impersonate the target in three different scenarios. The researchers did not actually use the information collected to commit fraud, but they showed that this would not have been difficult, even to get a temporary ID from the police.
3. "Applicability of lifestyle-routine activity theory to harassment in massively-multiplayer online role playing games" presents the results of a questionnaire about online harassment (N=222). The experiment is a repeat study of an experiment by Holt and Bossler [Hol09], focusing on a more specific setting.
4. "Efficient Drive-by-Download Detection" describes a tool that can detect whether a website has been infected with certain type of malware. The tool allows subjects to surf the most dangerous places on the net without having to worry about drive by downloads [Veg09]. Unfortunately, the students just managed to build the tool and were unable to allow subjects to use it.
5. "How dangerous is Geotagging?" describes a survey where subjects (N=22) where asked what they thought about the implications of geo-tagging photos for their privacy. The researchers found that the level of concern is low.
6. "Analysing malicious Tor exit traffic" presents an empirical study whereby traffic from a Tor exit node set up by the researchers was analysed to identify which countries are the most popular targets of attacks.
7. "Phishing using QR codes" describes a Randomised Control Trial (RCT) where phishing targets (N=57) were recruited by persuading University staff and students to use their smart phone to scan a QR code that was printed on 35 posters. The QR code led to a web site with a questionnaire about campus facilities. The experimental groups were served an anti-phishing warning, and the control group did not. The results indicate that QR codes are effective bait and that warnings do help but not enough.
8. "Influencing people's illegal downloading behaviours using warnings and other emotion-inducing visuals" describes an experiment whereby subjects (N=59) were shown 7 different types of warnings designed to make them think again before actually downloading content from a web site set up by the researchers. The results indicate that positive warnings (such as make sure that you don't put yourself at risk by committing an offense) were more effective than negative warnings (such as downloading is theft). The researchers did not include a control group.
9. "Creating phishing awareness in students" describes an RCT whereby 605 students were sent a phishing email to which N=66 subjects responded. The experimental group were sent several emails to warn them about phishing, the control group received no warnings.

From the control group 27 subjects entered PII into the phishing site, and 17 did not. This suggests that anti-phishing warnings might help a little.

10. “Understanding Users' behaviour towards online privacy” describes a pre-test/post-test experiment designed to test the ability of the subjects (N=67) to learn about privacy technology. All subjects were asked to complete a questionnaire about a certain privacy technology. 30 subjects volunteered for an information pack and 13 of those completed a second questionnaire to see what they had learned. The results are inconclusive, as it cannot be ruled out that only those subjects who already knew about the privacy technology took part in the post-test.
11. “Stimulating litter removal in community rooms through interactive trash cans” describes how an interactive waste bin could improve the tidiness of the subjects (convenience sample, N=24), as compare to a control group. Subjects could vote for things (e.g. Pepsi vs Coke) by throwing used plastic cups in the right bin. The results were inconclusive.
12. “The dark side of Facebook” investigates the proclivity of subjects (N=28) to become friends with unknown people represented by two fake profiles on Facebook (one male, one female). Males were more likely to accept invitations from an unknown female than vice versa.
13. “Awareness to Cyber-crime of Higher education students” describes an empirical study whereby USB sticks infected with a “friendly virus” were lost in public places in order to see what the subjects (N=40) who found the USB sticks would do. About half the subjects inserted the USB stick in their PC, which duly reported home this fact.