Teaching Engineering students to "Think thief"

Pieter Hartel and Marianne Junger, University of Twente

(Version 17, 26 April 2014)

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

Word count main body of the text: 3967

Abstract

We report on an educational experiment where information technology students were encouraged to think out of the box about the dark side of information technology. Instead of taking the usual point of view of the engineer we challenged the students to take the point of view of the motivated offender. After teaching the course three years, we report on the exciting ideas our students came up with, and on the lessons we learned in designing and teaching the course. The main conclusions are (a) thinking thief inspires students to design creative projects, (b) working with real subjects creates a powerful learning experience, and (c) students are struggling with methodological issues.

1. Introduction

How does an engineer avoid criminogenic designs? The deceptively simple answer is: by taking the point of view of a motivated offender, who is looking for ways in which a design can be misused. This idea has occurred in different communities. For example in the cyber security community this is called "adopting the hacker mind-set" [Con11], and in environmental criminology this is called "thinking thief" [Gam03]. A simple example of what thinking thief entails is the ordinary beer glass, which, when broken, becomes a dangerous weapon. But the designer who has been thinking thief uses laminated glass that does not break [McG10a].

How should one teach thinking thief? This is a difficult question to answer because there are many ways in which designs can be misused. Stimulating the creativity of the designers is essential here. We believe interacting with real adversarial subjects gives a powerful boost to the creativity of the students.

Curricula across the world offer a variety of courses on thinking thief; even at pre-college level [Gon13] but cyber security curricula often include this aspect [Bra07]. However, the majority of cyber security courses focus on the technical aspects of thinking thief, for example cryptanalysis (mathematics), kernel hacking (systems), and red-blue teaming (networks). From the 16 ideas mentioned by Conti et al [Con11], only one involves real subjects. Schneider [Sch13] discussed the disadvantages of thinking thief (he calls it adversarial thinking), chief of which is that in a technical context students get bogged down into idiosyncratic implementation details. In a social context, there are no implementation details, as the object of study is the human opponent, not the technology. In almost all courses that do take the human factor into account the students are either pitted against each other [Con12], or against a simulated opponent. There is a good reason not to engage the students with real opponents: the ethical and legal issues are thorny [End03].

We have been experimenting with a course in thinking thief where the students take on real opponents. For example, suppose that students wish to investigate whether shoppers in an online market place are able to distinguish fake from bona-fide adverts. Few Institutional Review Boards (IRB) would allow researchers to publish fake adverts on a public site. Yet our students found an ingenious solution that our IRB had no problems with. They found already existing fake adverts on the site and presented those to their subjects.

The purpose of this paper is to report on our experience teaching information technology students to think thief, presenting some of the most creative ideas of our students, and hopefully to inspire colleagues.

2. The Course

Our course is attended by master students from a Cyber Security program and bachelor students from a Creative technology program. The four key ingredients of the course are:

1. Social science research methods [Doo09]. During the first weeks of the course, the students were the researchers in a re-enactment of a classic experiment, such as Milgram's lost letter experiment [Mil65]. We used USB keys instead of letters. The students learnt about formulating hypotheses, designing an intervention, analysing the results, and interpreting the results in the light of a theoretical framework.
2. Legal and ethical responsibility [End03]. The students had to design and document their own experiments according to the requirements of the IRB.
3. A theoretical framework [Cla08]. We use the framework of situational crime prevention, and more specifically opportunity theory, which basically states that "opportunity makes the thief". Therefore effective crime prevention techniques should reduce the opportunity by standard means such as (1) increasing the effort, (2) increasing the risk, and (3) reducing the rewards for the offender. For more detail on studies into the effectiveness of opportunity reducing techniques visit the U.S. Department of Justice sponsored Center for Problem Oriented Policing[1].
4. Practical experience [see the online appendix]. The students had to study the literature on their chosen topic, write a research proposal, seek approval of the IRB, build the necessary web sites, tools, and services, recruit subjects, collect the data, and finally write and present a paper.

The course was taught in 2011, 2012 and 2013, when it was attended by a total of 108 students working in teams of 2 or 3. 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.

During the first few weeks the student teams drafted a research proposal. After two rounds of feed-back on the research proposals, the teams were given permission to submit their projects to the IRB. Upon approval of the IRB, the students went on to execute their projects. This resulted – over three years – in 43 six-page papers that were presented at a half-day conference at the end of each instance of the course. The students reviewed each other's papers and presentations, moderated by the lecturers.

[1] http://www.popcenter.org/

A project would either be a survey or an experiment. A survey is usually person oriented, and focuses on the differences between subjects. An experiment is, in general, 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 existing situations, e.g. do you accept all friend requests on face book? Experiments push thinking thief further than surveys because experiments assess behavioural change, e.g. to what extent does an anti-phishing warning really work?

Planning and executing an experiment is harder than a survey for a variety of reasons. For example, an experiment sometimes involves deception, and in that case the real purpose of the experiment can only be disclosed at the debriefing stage. This requires subjects to sign an informed consent, and requires debriefing of the subjects where possible. 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 a survey, that we encouraged our students to design experiments, and we accepted surveys somewhat reluctantly. Over the three years 20 teams performed a survey and 23 teams performed an experiment.

A quality experiment takes considerable time and skills to prepare and execute. The time requirements range from months to years and the ideal skill requirements include a bachelor degree in the social sciences. We were therefore faced with two problems. Firstly, our course is a one-semester course (6 European credits, 168 hours). 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, not all our students have a formal training in social science research methods. Hence we were faced with the challenge to develop a "lightweight" approach to social science research for information technology students.

3. Lessons learned about the IRB

During the first instance of the course there was no IRB in operation. Therefore, 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.

The second time we ran the course, the IRB had just been established. This brand new IRB was not always able to suggest workable alternatives when it considered an idea too risky. We discuss the three most important aspects of our experience with the IRB.

Firstly, the IRB had concerns about some of the more realistic Crime Science experiments. For example one team of students wanted to create a fake company with a convincing looking web site, Facebook page, and Twitter account. The objective was to assess how a convincing presence in all these areas would social engineer subjects more successfully than just a web site. The IRB considered the risks too high and could not suggest viable alternatives. In the end the team dropped out of the course out of frustration. Another team who wanted to do something similar with fake adverts did find a clever way forward, as we discussed in the introduction.

Secondly, the IRB was deeply concerned about the well-being of the subjects, which sometimes led to an unworkable solution. For example, several teams wanted to interview students in the library. The committee rejected this because it felt that this would disturb

the students in the library too much. Instead, the committee suggested that the teams should prepare leaflets with an invitation, which could then be deposited in the library, subject to permission from the library staff, for potential subjects to be picked up. One of the teams followed the suggestion and found themselves without any subjects. Other teams approached their subjects on the campus instead.

Thirdly, the form to be completed for the IRB was unnecessarily complex and contained redundant questions. For example a question on accidental medical discoveries is not relevant for Crime Science experiments. Some teams got confused and handed in incomplete and inconsistent forms, which frustrated both the committee and the students. In the end 4 out of 13 teams from the second instance of the course completely changed their project and 2 teams dropped out because of the frustrating experience with the IRB. Some sample project ideas that did not make it past the IRB were: (1) Set up a phishing web site and ask subjects whether they are using the same password as on other sites, (2) Sniff usernames and passwords from public networks to discover how often subjects reuse them, and (3) Set up an evil-twin of a wireless network access point to observe whether subjects notice the difference.

We feel that the IRB as it functioned during its first year of existence erred a little too much on the side of caution. An IRB in medical studies discuss issues of life and death in randomized controlled trials. The issues at hand in information technology are – fortunately – not about such weighty decisions. It is also the responsibility of the University that our students become well prepared defenders of the Internet.

4. Studies executed by the students

Table 1 summarises the surveys on the left and the experiments on the right. It also shows the number of subjects taking part in the study and the number of subjects in the control group for the intervention. The number of subjects was relatively low and many findings did not reach statistical significance. In some experiments the students were unable to include a control group, even though they had planned to do so in their research designs.

The students recruited subjects from two different populations. In 11 studies subjects were recruited from outside the university (e.g. via online fora, advertisements, or from the staff of a hospital, or people on the street). In the remaining studies subjects were friends, colleagues or family. From each population, students selected subjects in two different ways. In 3 studies the subjects were selected via random sampling. Convenience sampling was used in all other studies. All subjects were thus real subjects and sometimes behave unexpectedly. This put the student researchers under considerable stress and we believe that this contributed significantly to the intensity of the learning experience for our students.

Table 1 Summary of the essential statistics of the 43 student papers. N=total number of subjects. C=total number of subjects in the control group(s).

Surveys	N	Experiments	N	C
First instance of the course				
1. War driving	0	7. QR-code anti-phishing training	57	12
2. Simulated ID theft	100*	8. Illegal download warnings	59†	0

3. 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
Second instance of the course				
14.Open email attachment	120	19.Remember multi passwords	23†	9
15.Drawing passwords	77†	20.Is an SLL certificate valid	26	11
16.How up-to-date is Java	70	21.Same as #7		
17.Social sports sites leak ‡	308*†	22.Same as #7		
18.Prizes and privacy	74†	23.Can you identify fake adverts	30	20
		24.Explore unattended computer	52	0
Third instance of the course				
25.Exposing job scams	15	35.Avoiding fake login screens	60	26
26.Privacy settings on Facebook	48†	36.Stress and Distraction	50	25
27.Gathering Information an Social Proof	50*	37.Privacy sensitive data	72†	14
28. Password Habits and Background	135	38.Peer pressure	51	26
29. Password Characteristics	4†	39.Social Engineering Keys	44	31
30.Picture Passwords	62	40.Cooperative Behaviour	762	434
31.Man-in-the- middle on wireless access point	59	41.Alternative trust	20	10
32.Juice jacking	31	42.Phishing Recognition Skills	20	1
33.Awareness online privacy	108	43.Phishing Awareness for children ‡	149†	79
34.Android user permission awareness	31†			

*=Random sample; †=Subjects recruited from outside the university (e.g. via FaceBook, Google Adwords, Android Playstore); ‡= An article has been/ is being submitted to a journal

The appendix provides a short paragraph on each of the 43 student papers but here we should like to mention four interesting ideas that have not received much attention in the literature yet.

Paper #2 presents a simple and ingenious idea. 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 needed from the Internet. 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, with success, to the police to obtain a temporary ID.

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.

Paper #17 investigates how easy it is to discover the home address of subjects from their web presence. Thinking thief, the students hypothesized that people are more likely to "leave their tracks" on a social sports site if they feel proud of an achievement. The students collected the run keeper[2] profile of 304 subjects, and calculated the home address from the set of tracks of each runner. Since most people start running from home, and stop running to cool down close to home, the address could be determined accurately in most cases. The students then tried to obtain the home address also from other sources, such as the Face book profile of the runners. Discovering the home address from run keeper profiles was twice as successful as from Facebook. This work has been heavily revised by the staff and has been accepted by a journal.

Paper #43 researched the effect of anti-phishing training on 159 school children aged 9-12. There was a statistically significant difference between the experimental group, which received training and the control group, which did not receive training. One of the most interesting suggestions from this study is that children and adults react differently to anti-phishing training. This work is being revised by the staff and will be submitted to a journal for publication.

We found the creativity of our students combined with the research experience of the staff produces novel results that are worthy of publication in journals.

5. Lessons learned from the students

We asked the students who took the course in 2013 to complete a questionnaire[3] at the beginning of the course and to complete it once more at the end. Of the 50 students who attended the course in 2013, 64% completed the questionnaire when the course had just started and 98% completed the questionnaire at the end of the course. We then compared the results of the two sets of responses at the aggregate level. We could not combine the findings at the individual level because the questionnaires were anonymous.

At the pre-test and at the post-test we asked the students to rate on a five-point scale (1=Poor, 2=Fair, 3=Good, 4=Very Good, 5=Excellent) how they would assess their (a) level

[2] http://runkeeper.com/
[3] https://docs.google.com/forms/d/1yaAM1nI2nassZZKmQ6ZP_Pd0wgpTIKfzPwoUa_KwKVo/viewform

of knowledge of social science methods, (b) ability to use statistical software, and (c) ability to social engineer people.

Table 2 Assessing the differences on the dependent variables due to taking the course

	Pre-test N=32		Post-test N=49	
Dependent variables (five point scale)	Mean	Std dev	Mean	Std dev
(a) Level of knowledge of social science methods *	2.1	0.82	2.7	0.78
(b) Ability to use statistical software	1.8	0.78	2.1	0.89
(c) Ability to social engineer people	2.6	0.88	2.8	0.74

* $p < 0.02$

Table 2 shows the mean level of knowledge at the pre- and post-test on the five-point scale. The averages for the level of knowledge are significant, so we may conclude that students have increased their knowledge by 0.6 on the five-point scale. This represents a medium to large effect size [Coh92]. The remaining averages are not statistically significant, but they do show increases in the expected direction

We also asked the students to give their opinion on the course. Figure 1 shows that few students (80%) had sleepless nights (a). Most felt that they worked as hard in this course as for other courses, although 29% mentioned that they worked harder on the present course (b). 56% stated that the course increased their interest in social science research (c). 51% of the students were confident that, during their career, they would use what they learned during the CCS course (d). 65% agreed that they would never forget the CCS course (e). 59% I had a lot of fun doing the CCS experiment (f). 49% are confident that they are less likely to fall for a social engineering attack, as a result of taking the CCS course (g).

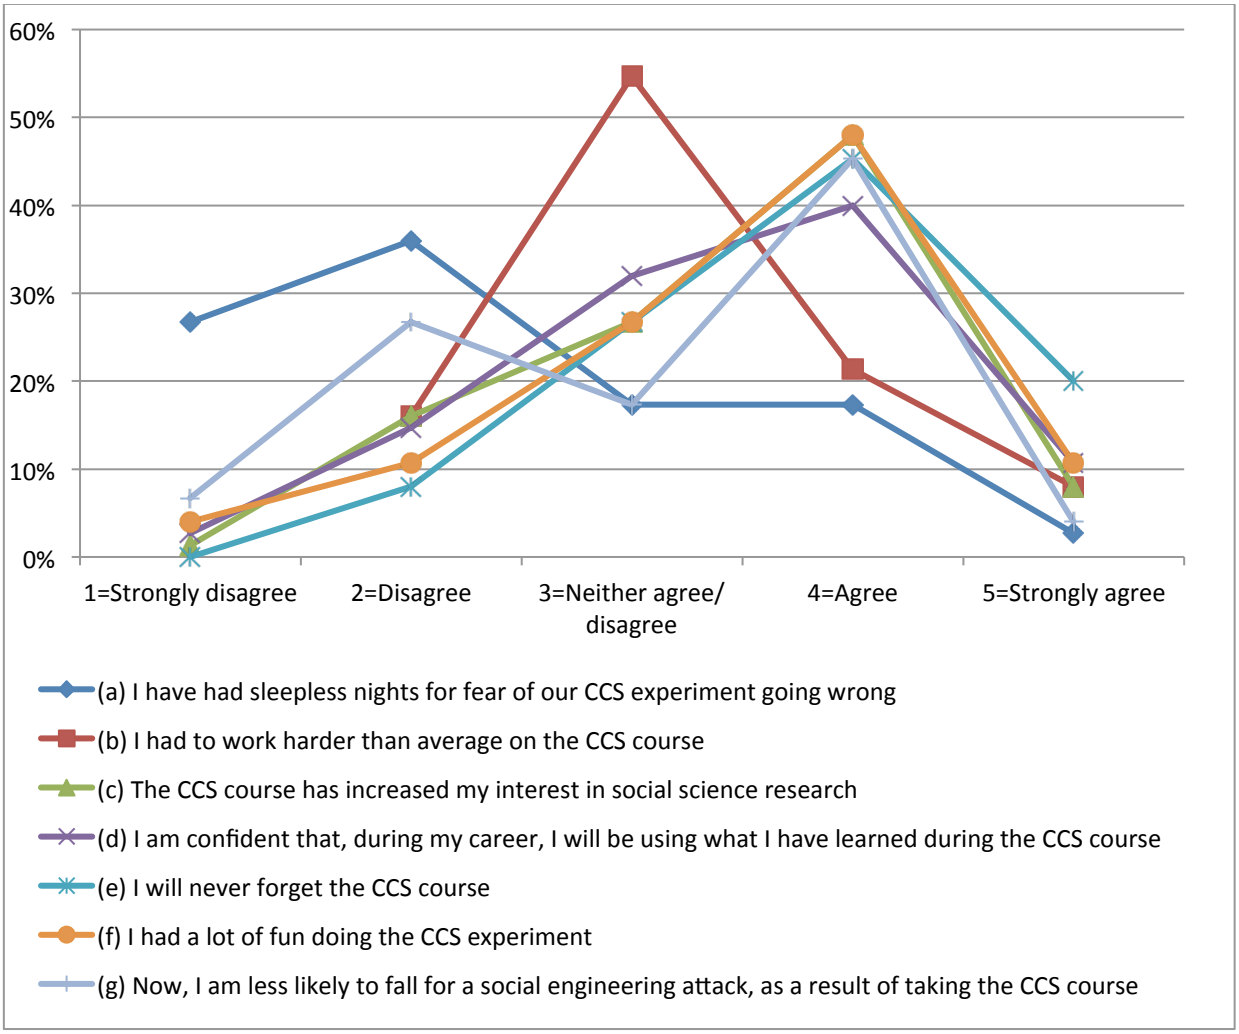

Figure 1 Opinions of the students on the course

We interpret these results as evidence that students found thinking thief challenging but exciting. We believe that is due to the fact that the students were free to choose their own topics on a subject that triggered their imagination, and due to the realistic nature of the experiments with real subjects. We also believe that the students appreciate the opportunity to execute a complete social science research project from start to finish, as a kind of mini MSc project.

The students had to work hard to be able to apply a range of skills, including experimental design, basic statistics and SPSS [Fie09], and the APA reporting standards [APA08].

Our course offers students an experience that is comparable to what medical students experience when they train in a hospital, but on a smaller scale. Where medical students work with real patients, our students work with real potential crime victims . The fact that this is for real is a powerful incentive to work hard.

6. Lessons learned about engineering curricula

Engineers sometimes make assumptions about misuse of technology that are not validated in the real world. For example, is a weak password an important crime facilitator? If not then why should we worry about weak passwords? Our course inspired students to consider not only the technology but also its consequences, which is what thinking thief is all about.

Only a few of the engineering curricula at our university offer a basic course in social science research methods. This means that not all engineers learn how to research the impact of their work on the end user. Our course addresses this gap, but it only touches upon the most essential aspects of social science research methodology.

Empirical research requires a mind-set that is different from an engineering mind-set. Engineers more often focus on logical arguments, and the effect of technology on the individual, whereas we believe that engineers should pay more attention to statistics and the effect of technology on the population at large.

7. Conclusions

Conducting a truly multidisciplinary course was challenging for the teachers as well as the students. There were two main issues that we were confronted with. Firstly, the Institutional Review Board needed to gain experience, and sometimes frustrated the students. Secondly, engineering students hardly had any skills that fitted with social science research model and therefore needed a lot of methodological guidance for conducting the experiments.

Setting aside the issue of the missing control groups, all 43 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 an answer to the challenge we set ourselves of developing a “light weight” approach to Crime Science: simply drop the requirement that the study is based on sufficiently large groups as one would expect for ‘normal’ scientific studies.

Our empirical study shows that students who completed the course have more knowledge of social science methods, and estimate their ability to use statistical software as well as their ability to social engineer people to be higher than students who did not follow the course, after controlling for possibly confounding variables.

The students who completed the post-test questionnaire gave the course positive evaluations. The majority mentioned that the course increased their interest in the social sciences, and that they would never forget this course and that it was fun to do.

These findings are limited by the fact that those who did not follow our course were not randomized but self-selected, and that self reported levels of knowledge are not objective.

Despite the limitations, these findings support the thesis that the CCS students benefitted significantly from following our course focusing on topics that we believe to be important.

We are still keen to improve our course. We hope that this article will help others with similar aims, and that we may hear from other teachers who have solved similar problems.

References

[APA08] APA Publications and Communications Board. Reporting standards for research in psychology: Why do we need them? what might they be? American Psychologist, 63(9):839-851, Dec 2008. http://dx.doi.org/10.1037/0003-066X.63.9.839

[Bat06] M. Bateson, D. Nettle, and G. Roberts. Cues of being watched enhance cooperation in a real-world setting. Biology Letters, 2(3):412-414, Sep 2006. http://dx.doi.org/10.1098/rsbl.2006.0509

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

[Bra07] S. Bratus. What hackers learn that the rest of us don't: Notes on hacker curriculum. IEEE Security & Privacy, 5(4):72-75, Jul 2007. http://dx.doi.org/10.1109/MSP.2007.101

[Cla08] R. V. Clarke. Situational crime prevention. In R. Wortley and L. Mazerolle, editors, Environmental Criminology and Crime Analysis, pages 178-194. Willan Publishing, London, Jun 2008. http://www.routledge.com/9781843922803

[Coh92] J. Cohen. A power primer. Psychological Bulletin, 112(1):155-159, Jul 1992. http://dx.doi.org/10.1037/0033-2909.112.1.155

[Con11] G. Conti, T. Babbitt, and J. Nelson. Hacking competitions and their untapped potential for security education. IEEE Security & Privacy, 9(3):56-59, May 2011. http://dx.doi.org/10.1109/MSP.2011.51

[Con12] G. Conti and J. Caroland. Embracing the Kobayashi Maru: Why you should teach your students to cheat. IEEE Security & Privacy, 9(4):48-51, Jul 2012. http://dx.doi.org/10.1109/MSP.2011.80

[Doo09] D. Dooley. Social Research Methods. Prentice Hall, fourth edition, May 2000. http://www.pearsonhighered.com/educator/product/Social-Research-Methods/9780139554285.page

[End03] B. Endicoytt-Popuvsky. Ethics and teaching information assurance. IEEE Security & Privacy, 1(4):65-67, Jul 2003. http://dx.doi.org/10.1109/MSECP.2003.1219073

[Fie09] A. Field. Discovering statistics using SPSS. Sage, London, 3rd edition, Jan 2009. http://www.uk.sagepub.com/field3e/main.htm

[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

[Gon13] M. Gondree, Z. N. J. Peterson, and T. Denning. Security through play. IEEE Security Privacy, 11(3):64-67, May 2013. http://dx.doi.org/10.1109/MSP.2013.69

[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

[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. http://dx.doi.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), page 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

[Mil65] S. Milgram, L. Mann, and S. Harter. The Lost-Letter technique: A tool of social research. The Public Opinion Quarterly, 29(3):437-438, 1965. http://www.jstor.org/stable/2746945

[Sch13] F. B. Schneider. Cybersecurity education in universities. IEEE Security & Privacy, 11(4):3-4, Jul 2013. http://dx.doi.org/10.1109/MSP.2013.84

[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

Online Appendix

The main idea and the results of the 43 student 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.
14. “Unknown files and personal relationships” presents a survey where subjects (N=120) completed a questionnaire about opening email attachments. Unsurprisingly, the closer the sender and the receiver are acquainted the more likely it is that that receiver will actually open the attachment.
15. “Viability and usability of drawing based password systems” presents a survey where subjects (N=77) created 1092 drawings representing passwords. The subjects we asked to re-draw their own passwords and those of others and they we asked to determine whether two passwords were the same. The results were inconclusive, partly due to the fact that drawing was done using a mouse.
16. “Survey on computer use and (un)safe habits with JRE” presents a survey where subjects (N=70) were interviewed about the state of the Java implementation on their laptop. Only 23% of the subjects used an up to date system. The only significant result the researchers found is unsurprising: there is correlation between checking regularly for a new version and actually having a recent vision installed.
17. “Mining social sports: Leaking private information through social sport networks” presents a survey where two different attempts were made to determine the home address of the subjects (N=308). The home address of 132 subjects could be determined from the run keeper profile of the subject, whereas other sources on the Internet led to the home address of only 64 subjects. The difference is significant, and shows that social sport network users are not sufficiently aware that their privacy is at risk.
18. “Exploring the effect of external factors on disclosure of private information” presents a survey where subjects (N=74) were asked to disclose personal information. The questionnaires for the control and experimental groups differed only in the order of the questions and whether the subjects could win a prize. There were no significant differences between the control and experimental groups, but offering a prize seems to reduce the number of subjects that provide private information.
19. “Multiple password simplicity” presents an experiment where subjects (N=23) were asked to choose a password and to reproduce it 4 days later. The control group of 9

subjects was asked to choose a standard password and the experimental group of 14 subjects was asked to choose a multi word password. There were no significant differences between the control and experimental groups.

20. “User performance in making trust decisions concerning SSL encrypted connections” presents an experiment where subjects (N=26) were asked to identify valid SSL certificates. The experimental group of 15 subjects received some training using a specially designed game, whereas the control group of 11 subjects received no training. There were no significant differences between the control and experimental groups.
21. Same as #7
22. Same as #7
23. “Awareness of scam advertisements within C2C auction sites” presents an experiment where subjects (N=30) were asked to identify fake adverts on an online auction site. Two experimental groups (of 5 subjects each) were shown different videos to warn against fake ads; the control group (20 subjects) was not shown videos. Then all subjects were tested again, with a new set of fake adverts. There were no significant differences between the control and experimental groups.
24. “The abuse of digital data” presents an experiment where subjects (N=52) were exposed to unattended computers showing the site of an online social network or online bank. Different signs warning against unauthorised use of computers were placed near the computers. There was no control group. Only 4 subjects actually spent a few minutes exploring the sites on offer.
25. “Exposing job scams” presents a survey where marketing and computer science students (N=30) are asked to identify job adverts on Craigslist that are too good to be true. Each subject was presented with the same 30 adverts, 15 of which were fake and 15 real. About 75% of the adverts were classified correctly. None of the independent variables, such as degree course or sex of the subjects made no significant difference to the classification accuracy.
26. “Privacy versus standard settings regarding Facebook” presents a survey on the privacy awareness of Facebook users (N=48) that is basically a repeat of danah boyd’s research [http://firstmonday.org/ojs/index.php/fm/article/view/3086/2589]. None of the results are significant.
27. “Gathering Information With Or Without Social Proof Through E-mail And Telephone” was an experiment designed to compare the effectiveness of phishing via email and via the phone. However, the management of the university terminated the experiment as soon as the N=50 subjects had been first sent a phishing email. The unfortunate researchers had chosen a topic for their phishing emails that management felt might jeopardize one of their policies. Neither the ethical committee, nor the lecturers had foreseen this eventuality. The students were marked on the basis of a short note on their experince.
28. “Secure Password Habits Depend On The User's Background” reports on a survey of 135 snowball sampled subjects that were asked how they choose and managed their passwords. Respondents with an ICT background do marginally better than other respondents. No statistically significant results were found.
29. “Examining Password Characteristics using Dictionary Attacks and Structured Qualitative Interviews” reports on four structured interviews on password policy with industry professionals.
30. “Gesture Frequency in Picture Passwords search for the weak passwords of Windows 8 picture password” reports on a study where 62 university students were asked to draw a picture password twice. Only 52% remembered their password after two weeks. The

main result was that picture passwords do not seem to offer obvious default choices that textual passwords suffer from.

31. “Removing SSL using man in the middle on a wireless access point” investigates security awareness. 25 students from the social sciences logged in to a rogue wireless access point, and 34 science students logged in to another rogue wireless access point. The researchers had put up banners and distributed flyers near each rogue access point to warn prospective users not to disclose their credentials to unfamiliar access points. 44% of both communities ignored the warnings. No statistically significant differences were found between students from the different faculties, nor were the effects of the banners or the flyers found to be statistically significant
32. “Understanding victims of juice jacking” investigates whether university students trust a charger for smart phones in a public place. The researchers had built a charger that could, without the owner’s knowledge, in principle infect any device connected to via the USB port. From the 31 potential subjects only 12 connected their smartphone to the charger; all subjects were debriefed and interviewed. No statistically significant results were reported.
33. “Raising awareness over online privacy on social networks” presents an online survey, asking 108 Facebook users about their privacy awareness. The only statistically significant result is that male subjects are more aware of privacy issues than female participants.
34. “Android users' awareness of giving permission to new applications” reports on a survey where the researchers developed a Sudoku App and published it on Google Market. 31 subjects downloaded the App but only one user actually noticed that the App requested more permissions than it should have.
35. “Avoiding fake login screens” studies whether university students remember to press CTRL+ALT+DEL before entering their user name and password on a specially prepared, apparently unused PC (N=60). The experimental group of 34 subjects is confronted with a poster next to the PC with an appropriate warning. The poster has a small beneficial effect that is not statistically significant.
36. “Experimenting With Stress and Distraction Factors Alongside Phishing E-mails” asks 50 university students to decide which of 30 emails are phishing emails. The experimental group is put under stress by limiting the time per email, and by displaying unrelated pop-ups to the subjects. The control group has no time limit and is not distracted by pop-ups. The experimental group performed twice as bad as the control group. Whether this result is significant or not has not been reported.
37. “On testing the effectiveness of methods to prevent the relinquishment of privacy sensitive data” presents a number of interventions designed to warn subjects against giving away private information via Facebook connect. The researchers recruited 72 subjects via Google adwords to a landing page where subjects could vote for, or against a political statement. There were two types of warnings, but neither had a significant effect.
38. “Using peer pressure for phishing on social networks” reports on an experiment (N=51) with two groups of university students who are also Facebook users, and who are confronted with a phishing link. The experimental group is made to believe by positive comments from the researchers that the link is bona fide. The researchers did not interact with the control group of 26 subjects. No statistically significant effects of the interaction have been observed.

39. “Social Engineering Key Experiment” is a repeat study of an experiment with an intervention that warns university staff and PhD students against handing over their office keys to strangers. The intervention was found to be effective.
40. “Analysis on Cooperative Behavior While Being Watched” is a repeat study of Bateson et al’s “Cues of being watched” [Bat06] in a cafeteria run for and by university students. Similar to other researchers who tried to replicate Bateson et al’s results this research found that pictures of eyes do not necessarily improve cooperation.
41. “Alternative to building a trustful relationship when phishing” reports on an experiment where 20 pairs of Facebook friends on a university campus are pitted against each other. The experimental group was offered a chocolate bar to reveal profile information about their Facebook friend; the control group did not receive a bribe. No statistically significant differences between the two groups were found.
42. “Determining the Difference of Perception and Reality in Phishing Recognition Skills in Different Situations” studies the ability of university students to recognize phishing emails when put under stress. The experimental group was limited in time, whereas the control group was not. No statistically significant differences between the two groups were found.
43. “Phishing Awareness for children” researched the effect of anti-phishing training on 159 school children aged 9-12. There was a statistically significant difference between the experimental group, which received training and the control group, which did not receive training.