Technology in Society 63 (2020) 101382

Contents lists available at ScienceDirect

Technology in Society

journal homepage: http://www.elsevier.com/locate/techsoc

Ethics in cybersecurity research and practice

Kevin Macnish *, Jeroen van der Ham

Department of Philosophy, Department of Computer Science, University of Twente, Drienerlolaan 5 7552NB, Enschede, Netherlands

ARTICLE INFO

Keywords:
Cybersecurity
Ethics
Privacy
IRB
Personal data
Discrimination
Trust
Research

ABSTRACT

This paper critiques existing governance in cyber-security ethics through providing an overview of some of the ethical issues facing researchers in the cybersecurity community and highlighting shortfalls in governance practice. We separate these issues into those facing the academic research community and those facing the (corporate) practitioner community, drawing on two case studies. While there is overlap between these communities, there are also stark differences. Academic researchers can often rely on research ethics boards (REBs) to provide ethical oversight and governance which are typically unavailable to the practitioner community. However, we argue that even within the academic community the constitution of REBs is such that they may be (and in some cases at least are) unable to offer sound advice. Our recommendations are that ethics should be taught in far greater depth on computer science courses than is currently the case, and that codes of conduct should be developed and deployed provided they can be seen to be effective. In tandem with these, an active discussion regarding the ethics of cybersecurity and cybersecurity research is urgently needed.

1. Introduction

In this paper we argue that current methods of ethical oversight regarding cyber-security ethics are inadequate. These methods fail in at least two areas: university-based development and in the broader community of practising cybersecurity experts. In the former the problems stem from a lack of awareness among members of the computer security community and ethical review committees as to the nature of the ethical problems regarding cybersecurity. In the latter the problems are widely known, but a lack of adequate guidance or accountability forms a barrier to consistent ethical practice. We are not claiming that current cybersecurity development or practice are unethical. Rather, our point is that these practices go largely ungoverned and unguided, despite the clear potential for significant harm. We argue that there hence needs to be a greater appreciation of the risks of cybersecurity development in ethical review committees and clear codes of conduct for the professional community which cover both development and practice.

The paper opens with a case study regarding academic research into cybersecurity which was ethically flawed, but which genuinely sought ethics committee approval. This approval was denied, not because of the flaws in the case but rather because the case did not raise obvious issues of human subject research or personally identifiable information. This suggests that the ethics committees in the institutions consulted had a worryingly narrow view of ethical issues in their own field of research. We then list ethical issues which include, but go much further than, privacy and the confidential handling of personally identifiable information.

In the second part of the paper we look at a case study concerning research in the non-academic practitioner sector. Here again we note ethical flaws in the research which, in this case, arguably went unnoticed due to the absence of adequate ethical oversight. As with Part I, we follow the case study with a list of perceived ethical issues pertaining to practitioner research in cybersecurity. Some, but not all, of these issues overlap with those faced by the academic community. We conclude with a call for a mature discussion on ethical issues in the realms of cybersecurity research which embraces but also goes beyond concerns with privacy.[1]

The challenge facing governance, we argue, is that existing structures in university and other formal research environments can be insufficiently flexible in recognizing the ethical issues raised herein. This is illustrated by the first case study in which several Research Ethics Boards (REB) failed to protect the interests of research subjects in what,

* Corresponding author.
E-mail addresses: k.macnish@utwente.n (K. Macnish), j.vanderham@utwente.nl (J. van der Ham).

[1] As is standard for applied ethics papers, we do not explicitly draw on any one ethical tradition [84]; p. 3). While the principles discussed here are broadly consistent with most deontological, rule utilitarian, and intutionist frameworks, we see ourselves as operating in a Rossean tradition of highlighting *prima facie* duties which may at times conflict [85]. In such cases of conflict we would adopt a Rawlsian process of reflective equilibrium [86] to determine the preferred outcome.

https://doi.org/10.1016/j.techsoc.2020.101382
Received 3 December 2019; Received in revised form 2 September 2020; Accepted 4 September 2020
Available online 9 September 2020

we argue, was obviously ethically questionable research. On the other hand, where research takes place outside of these environments, typically in the private sector, there are no governance structures in place, leaving researchers with an even harder task when seeking ethical input to guide their research.

Finally, this paper is not intended to be a systematic analysis of ethical issues arising in cybersecurity. It is rather an ethical analysis of two case studies, combined with reflections from the authors' collective experience in teaching the subject over twenty years and additional research. An exhaustive exploration of all the ethical issues in cybersecurity goes well beyond the scope of a single research paper. Likewise, it is intended as a broad critique of the state of governance currently available to the cybersecurity community. It is not intended to be aimed solely at the nature of REBs. Indeed, the international variety of approaches to ethics review would require a systematic empirical review of REB constitution and practices. Nonetheless, to emphasise the global nature of our concerns, we have referred throughout to REBs rather than the US nomenclature of Institutional Research Boards (IRBs), or the European Research Ethics Committees (RECs) or other, except where these appear in quotes. Lastly, it is our hope that this analysis will raise awareness of ethical issues in cybersecurity which go beyond those most commonly acknowledged (such as privacy) and stimulate debate as to the state of governance in the cybersecurity community both within academia and without.

2. Part I - cybersecurity development in academic contexts

2.1. Case study 1: Encore

"Statement from the SIGCOMM 2015Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PCs view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted" [1].

The above warning was placed at the head of an article accepted by the SIGCOMM Program Committee in 2015. The paper in question concerned creating scripts to monitor levels of censorship. The scripts (called Encore) were then placed on the webservers of obliging companies (or on dummy advertising sites) and seamlessly transferred to the computers of clients when they visited those webservers. From clients' computers, Encore would then try to access sites that were likely to be censored and send information about their success or lack thereof back to the designers of the script. At no point were clients aware that Encore was running on their computer, still less were they asked for consent to have it operating on their computer.

At the time of writing, at least 17 companies had deployed Encore on their webservers, which led to "141,626 measurements from 88,260 distinct IPs in 170 countries, with China, India, the United Kingdom, and Brazil reporting at least 1000 measurements, and more than 100 measurements from Egypt, South Korea, Iran, Pakistan, Turkey, and Saudi Arabia. These countries practice some form of Web filtering" [1]; p. 662).

This last sentence is something of an understatement. In many of these countries, the mere act of visiting a banned site may lead to further investigations by the security services and may highlight individuals as persons of potential interest. This would be bad enough for a typical unwitting user, but if the client happens to be a dissident writing for free speech in their country, then the act of running the script from their computer could alert the security services to their activities. In the words of one of the SIGCOMM Program Committee members, the requests "could potentially result in severe harm: for example, when the user lives in a regime where due process for those seen as requesting censored content may not exist" [2].

As noted in the Statement from the SIGCOMM 2015 Program Committee (above), the authors did recognize some of the ethical concerns with their work. They determined only to have the scripts attempt to connect with sites that were not overly contentious, such as Facebook, YouTube and Twitter, and explicitly recognized that the research raised certain risks that were not fully understood [1]; p. 663). They go on to say that achieving balance between the risk to research subjects and the benefit of the research is difficult, but that, "striking this balance between benefit and risk raises ethical questions that researchers in computer science rarely face and that conventional ethical standards do not address" [1]; p. 663).

The paper proceeds to list attempts by the authors to have the measurement collection reviewed by REBs at two leading US universities [1]; p. 662). Somewhat surprisingly, given the in principle plausibility of gaining informed consent from research subjects,[2] both REBs declined to formally review the proposal as it did not "collect or analyse Personally Identifiable Information (PII) and [was] not human subjects research" [1]; p. 664). This, as we say, is surprising as human subject research is not merely that which collects or analyses PII, as the doctors at the Nuremburg Trials, upon whose direction much of contemporary research ethics is founded, would have been ready to point out.

The authors list several reasons for not requesting informed consent. These include the fact that "there are classes of experiments that can still be conducted ethically without [informed consent], such as when obtaining consent is either prohibitive or impractical and there is little appreciable risk of harm to the subject" [1]; p. 664). While this is true, this study does not appear to be a case in which getting informed consent would be prohibitive and, even if it were, there is clearly appreciable risk to participants and so the research would not class as low-risk observation. Indeed, as Byers notes, "PC members and survey respondents of an independent study agreed that most users for whom censorship is an issue would be unlikely to consent to Encore's measurements" [2].

A second reason given for deciding not to request informed consent was that doing so "would require apprising a user about nuanced technical concepts … and doing so across language barriers" [1]; p. 664). The researchers were concerned that "such burdens would dramatically reduce the scale and scope of measurements, relegating us to the already extremely dangerous status quo of activists and researchers who put themselves into harm's way to study censorship" [1]; p. 664). The desire to move beyond research that puts the researcher in harm's way to study censorship is well-motivated. However, informing a research subject about the potential harms of complex research, and doing so across language barriers, is standard practice for many researchers in the medical and social science fields. It is not clear, therefore, why this case should be treated differently.

The authors note further that, "informed consent does not ever decrease risk to users; it only alleviates researchers from some responsibility for that risk and may even increase risk to users by removing

[2] It would not have been unfeasible to request voluntary, consenting participation from research subjects in the countries under consideration. This may have been difficult given the desired scale of the research, but neither impossible nor undesirable.

any traces of plausible deniability" [1]; p. 664). This is cynical in the extreme and reads (to us) as post-hoc justification. It is true that informed consent does not decrease risk to research participants, but the point is that participants should be given the opportunity to decide for themselves whether they wish to take those risks, and not have those risks imposed by researchers. It is the duty of the researcher to describe the risks to the participants in such a way that the participants can make an adequate decision in accepting those risks. Indeed, this does "alleviate researchers from some responsibility", but does so in a controlled context.

Consent is a central aspect of post-war research ethics and has underpinned the Nuremberg Trials, the Helsinki Declaration, and virtually all subsequent writings on research ethics [3–10]. In its place, the authors write that, "we believe researchers should instead focus on reducing risk to uninformed users …. It is generally accepted that users already have little control over or knowledge of much of the traffic that their Web browsers and devices generate (a point raised by Princeton's office of research integrity and assurance), which already gives users reasonable cover. By analogy, the prevalence of malware and third-party trackers itself lends credibility to the argument that a user cannot reasonably control the traffic that their devices send" [1]; pp. 664–65).

We agree that researchers should focus on reducing risk to participants, but this should not come as a zero-sum game with informed consent. Both should be present. Finally, the point raised by Princeton's office of research integrity and assurance is worrying, not least because of its source. It may be that in a country in which censorship is rarely practiced such a defence might be plausible, but history has shown that security services in totalitarian states tend to be extremely sensitive to such activities and often prefer an overly cautious perspective that ends in innocent people being incarcerated.

Our arguments may seem harsh here, especially as the authors worked to discuss the research with "ethics experts at the Oxford Internet Institute, the Berkman Center, and Citizen Lab" as well as "the organizers of the SIGCOMM NS Ethics workshop, which we helped solicit, to ensure that its attendees will gain experience applying principled ethical frameworks to networking and systems research, a process we hope will result in more informed and grounded discussions of ethics in our community" [1]; p. 664). We applaud the efforts of the authors and share their hopes in more informed and grounded discussions of ethics in the community, and it is to this latter end that we focus this paper.[3]

We do not write this to condemn the authors or the REBs that allowed these experiments to proceed. However, the forgoing case study amply demonstrates the paucity of ethical awareness within the academic computer science community at both researcher and REB level.

2.2. Ethical issues arising in academic contexts

As noted in the introduction, developments in cybersecurity methodology, tactics and techniques occur at both the level of academic research and in research at a corporate and government (i.e. practitioner) level. This is not to say that academic institutions do not practice cybersecurity: they do. However, at the stage of practice, the academic institution becomes indistinguishable for the purpose of our argument from the corporate or government practice of cybersecurity. While many of these ethical issues will invariably overlap, each of these also raises its own concerns. Below, we present a summary of those issues experienced in our own (academic) work, combined with insights from the research findings of others (see, for instance, Ref. [11–13]. In Part II we list ethical issues pertaining to practitioner-led research. In neither case are we claiming to be exhaustive in our lists of ethical issues. However, as evidenced by.

Case Study 1, at least some of these issues are not obvious and hence awareness about their potential for harm should be raised.

We have structured the list of ethical issues in parts I and II according to the Menlo principles, published in 2012 to lend coherence to ethical oversight of cybersecurity research [14]. The fact that the Encore case (above) happened at leading research institutions three years after the publishing of the Menlo Report suggests that take-up of those principles has been slow at best. The Menlo Report offers four principles for ethical research (respect for persons, beneficence, justice and respect for law and public interest), following in the tradition of the Belmont Report, a key document for establishing research ethical principles in the US [15]; see also [16]. However, as will be seen, several of the issues which we have experienced in our own work as university researchers in ethics and cybersecurity, do not easily fit within the Menlo framework.

2.2.1. Respect for persons

"Participation as a research subject is voluntary, and follows from informed consent; Treat individuals as autonomous agents and respect their right to determine their own best interests; Respect individuals who are not targets of research yet are impacted; Individuals with diminished autonomy, who are incapable of deciding for themselves, are entitled to protection" [14]; p. 5).

2.2.1.1. Informed consent. Informed consent is one of the mainstays of research ethics [6]. The ability to and act of gaining informed consent from those who are affected by research stems back to the Nuremberg Declaration and is at the heart of the Helsinki Declaration and the Belmont Report [15,17–20]. This is seen starkly in the above case of testing censorship systems. It may also be a factor when the system is neither owned nor operated by the researcher. In such cases, should permission be required for the system to be tested?

The justification for informed consent is disputed as to whether it is rooted in the autonomy of the research subject [3,5,21] or in the principle of minimizing harm to the research subject [8]. However, whichever approach is correct, the seeking and gaining of informed consent has been the backbone of research ethics in which people may be harmed throughout the post-war period and cannot be lightly ignored.

2.2.2. Beneficence

"Do not harm; Maximize probable benefits and minimize probable harms; Systematically assess both risk of harm and benefit" [14]; p. 5).

2.2.2.1. Protection of subjects from inadvertent harm. It is wholly plausible that there are cybersecurity research projects in which people may stand to suffer as a result of that research, again as illustrated in the Encore case. We take it as given that harm is not intended on research subjects, but an absence of intention does not amount to an absence of effect: unintended harm is harm nonetheless. However, there may be some confusion here as to standard ethical practice. In arguing that cybersecurity research ethics should draw from clinical research ethics, Tyler Moore and Richard Clayton, for instance, argue that in the event of recognizing harm arising, researchers should only stop the trial when the results are "statistically significant and the divergence in treatment outcome is substantial" [22]; p. 15). However, this is not standard practice. There is an additional principle of minimization of harm to the participant which overrides the interest of the research, particularly in cases where little or no consent has been given.

2.2.2.2. Privacy. Conducting research will often reveal personal data/personally identifiable information (PII), which then needs to be handled appropriately. There is a vast body of work regarding the definition, scope and value of privacy which directly pertains to research ethics [23]. However, any detailed discussion of a single ethical point will quickly extend beyond the scope of this paper. Furthermore, despite this volume of discussion, privacy issues continue to arise in cybersecurity research (see, for example, [24–28].

[3] For further, more detailed ethical discussion of the Encore project, see Ref. [87].

Legal sensitivity to privacy concerns is markedly varied, with the European General Data Protection Regulation [29] imposing strong restrictions and punitive measures to any engaging with data pertaining to European citizens, while the data of individuals in the Americas, Asia or Africa are far less subject to regulation. This could lead to "data dumping" in which research is carried out in countries with lower barriers for use of personal data rather than jump through bureaucratic hurdles in Europe. The result is that the data of non-European citizens are placed at higher risk than that of Europeans.

2.2.2.3. Reporting incidental findings. In the course of discovering personal data/PII, further information relating to an individual or organisation may be discovered [13]. Decisions need to be made in advance as to whether and how to inform that entity if appropriate. For example, evidence may emerge that a member of an organisation is seeking employment elsewhere, or that the spouse of an employee is having an affair with another employee. In the absence of a policy written in advance, such discoveries become ethical dilemmas in a way which they need not be. We are not aware of any academic research which has looked at the need to have a policy on incidental findings arising through cybersecurity research. Just as the Menlo Principles drew on the findings of the Belmont Report, though, it would reasonable to begin this process by drawing on the experience of the medical profession in dealing with incidental findings arising through examinations and clinical trials [30–35].

2.2.2.4. Testing the security of the system. Drawing back to the Encore case above, there is a question regarding whether leaked vulnerabilities should be used to install code on systems. In some cases, the researcher must act like a malicious party in order to fully test the system, and this will include using vulnerabilities. There are similar instances in which the researcher may want to engage in phishing tests, acting like a malicious agent in the choice of methods used, in order to determine vulnerabilities. Yet phishing by its nature employs deceit and one cannot easily gain prior informed consent from research subjects for fear of compromising the research [36]. On small-scale, limited participation experiments researching without prior informed consent and/or using deceit when the harms are minimal is typically take to be acceptable practice, as, for example, when engaged in some psychological research in which the aspect under investigation is other than that which the participants believe to be the case. However, harms are more difficult to predict, and the lack of consent more problematic, when the experiment extends beyond the scope of a few research subjects. Phishing and the use of vulnerabilities to test a system could cause extensive harm to those involved, which is exacerbated when no informed consent has been obtained. All harms should be avoided if at all possible. Where avoidance is not possible ethics committees can be beneficial in helping determining the proportionality of the harm to the research.

2.2.3. Justice

"Each person deserves equal consideration in how to be treated, and the benefits of research should be fairly distributed according to individual need, effort, societal contribution, and merit; Selection of subjects should be fair, and burdens should be allocated equitably across impacted subjects" [14]; p. 5).

2.2.3.1. Bias. The Encore project provides an example of how cybersecurity research can experience bias. By its very nature of attempting to determine the functioning of censorship firewalls, the research focused on users who lived, for the most part, in repressive regimes. Given the harms discussed above that were inherent for those users in the Encore project, it thus disadvantaged those already living in disadvantaged circumstances. A significant amount of cybersecurity research is carried out by researchers in the West who may have little to know experience of less advantaged groups living elsewhere in the world, which can, as in the case of Encore, lead to inadvertent bias against those groups. There has been a considerable body of recent work on bias in automated systems, much of which may have parallels in the cybersecurity realm where it may erroneously seem possible to isolate the individual affected from the system researched [37–39].

2.2.4. Respect for law and public interest

"Engage in legal due diligence; Be transparent in methods and results; Be accountable for actions" [14]; p. 5).

2.2.4.1. Coordinated vulnerability disclosure. Where vulnerabilities are discovered, should these be disclosed to a pertinent authority? Such an authority may be a company using the software which has the vulnerability, a third-party provider of that software, or a state entity which oversees vulnerabilities. In principle, a broad awareness of vulnerabilities is a positive as it can help the community come together to get a clear picture of how widespread the vulnerability is, whether any proprietary patches have been developed, and whether the vulnerability has been exploited. However, there is also the risk in broadcasting the vulnerability, even within a small community of cybersecurity professionals, that knowledge of that vulnerability will leak and could thereby be exploited. We will return to vulnerability disclosure in Part II, where it forms a significant part of Case Study 2. However, despite the literature on the value of and need for vulnerability disclosure (see, for example [40–46], we are aware of only a handful of university REBs which have a policy regarding vulnerability disclosure. Possibly even more than incidental findings, vulnerabilities are likely to be discovered in the course of cybersecurity research, and it is essential that those overseeing that research have clear guidance as to what should happen in those circumstances.

A further benefit of a vulnerabilities disclosure policy would be to protect the researcher in cases (such as Case Study 2, below) where vulnerabilities are discovered but no informed consent was obtained, and potentially was not obtainable. In such cases there is a high risk that the affected party will prosecute the university or researcher. In such cases, is there still a duty to make those discoveries known? What degree of risk should the researcher and the research institution each burden in investigating such vulnerabilities? The answers to will vary depending on the institution, but clarity is again essential to protect the researcher.

2.2.4.2. Testing on live and sensitive systems. Some systems cannot be taken off line in order to carry out research on them. This may be because they fulfil a vital function related to critical national infrastructure or because there is no built-in redundancy to the system [13]. In such cases, there is a risk of carrying out research that may have an impact on the functioning of that system. At the same time, such systems need to be tested for security purposes, possibly more so than their commercial counterparts. The preferable solution here would be for redundancy to be built into the system such that it could be tested a part at a time without risk to the whole, but this is clearly not always feasible. When this redundancy is not present and there is a risk of damaging the system, though, it is not clear how far the researcher should go in testing that system.

2.2.4.3. Impact on the commercial viability of a system. If vulnerabilities are found and not patched immediately, this could have an impact on the commercial viability of the system [13]. Does the researcher have a (whistle-blowing) duty to make such unpatched vulnerabilities public in order that greater pressure is put on the owner of the system to resolve the fault? Again, this is illustrated below in Case Study 2, but it is a problem which is faced by universities as well as commercial testers.

2.3. Recognizing ethical problems

While the authors of the Encore research did recognize and attempt

to seek assistance from more than one REB, this is not always the case. In many instances, researchers do not even recognize the potential for an ethical issue to arise. In discussing ethically-questionable research on the Tor network carried out in 2008, Christopher Soghoian notes that the researchers "simply did not see the ethical or legal issues associated with their data gathering" [47]; p. 146), although he goes on to quote one of the researchers as saying that they had been "advised that [seeing REB guidance or approval] wasn't necessary," suggesting that they had at least started to investigate the possibility of ethical issues arising [47]; p. 147). Following the presentation of the research in 2008, the University of Colorado announced that the researchers had not violated university ethics policies as "by any reasonable standard, the work in question was not classifiable as human subject research" (quoted in Ref. [47]; p. 148).

2.4. Competence of REBs

Ultimately, one of the key concerns of this paper is that REBs tend to consist of experts in ethics rather than experts in computer science, or vice versa. In our experience, which is echoed in the above case study, it is difficult to find an REB which effectively combines both sets of expertise. Nor are we the first to point this out, a similar point has been made repeatedly over the last decade: in 2008 [48,49], in 2009 [50], in 2010 [26] and in 2012 [11]; pp. 138–39; [12].

Despite being raised numerous times, the problem persists as many computer science researchers have received only elementary education in ethics, some of which might have included research ethics, while many ethicists have very little understanding of computer science research methods. This is not to say that there is no cross-over as there clearly is, but this is not as wide-spread as it needs to be in order to effectively oversee the developments with which we are concerned in this paper. Indeed, in research encompassing 700 REBs, Buchanan and Ess found that "in many cases … [REBs] did not know exactly what issues to consider as problematic or potentially harmful. IP addresses, clouds, worms, and bots are not part of the standard vocabulary of human subjects' research protections. For example, one respondent commented that "most REB members don't have degrees in [Computer Science]" [51]; see also [50]. When looking at 115 computer science courses which did contain ethics education, Fiesler, Garrett and Beard noted that a mere 19 specifically addressed cybersecurity [52].

The result is that, as noted in the above case study, REBs in computer science tend to react to well-recognized ethical and legal problems such as privacy and related issues regarding personally identifiable information. Less concern is directed towards the potential harm that may arise to individual research participants, particularly when they have not given or are unable to give informed consent to participate in the research. Standardly this is only permissible in cases of observational research in which the risk of harm is deemed (by an independent REB) to be low. However, cybersecurity research is often more interactive than mere observational research and the potential for harm may be considerable.

2.5. Summary

There are several ethical issues regarding university-based cybersecurity research, many of which we have highlighted above. While these are obviously of ethical concern, in many cases, such as that highlighted in Case Study 1, university REBs are simply not up to the requirements of offering effective oversight and guidance to researchers. This is primarily owing to the lack of joint expertise in computer science and ethics, which in turn stems from a weak commitment of computer science and philosophy departments at undergraduate level to teach ethics to computer scientists. While we accept that this is not a universal condemnation (apart from anything, the authors have each been teaching computer science ethics at universities for ten years, and the success of text books such as *A Gift of Fire* [53] attests that we are not alone in this), we are concerned that for significant institutions to miss the ethical problems in the Encore case suggests that the impact of this teaching has yet to filter through to research ethics oversight. Furthermore, as indicated above, there are several ethical issues pertaining to cybersecurity research (vulnerability disclosure, incidental findings specifically in cybersecurity) which do not standardly form a part of any university ethics policy and yet risk being encountered on a wide basis.

3. Part II - cybersecurity development in industry contexts

3.1. Case study 2 - MedSec

In August 2016, independent security research group MedSec purchased and attempted to attack a number of St. Jude Medical devices, including pacemakers and heart monitoring devices designed for home use. The team claimed to find multiple vulnerabilities in the home monitoring devices, including those which could be used to influence the behaviour of the pacemakers.

Rather than disclosing this to St. Jude Medical directly, MedSec teamed up with the investment firm Muddy Waters to short the stock of St. Jude Medical. They then released partial information about the vulnerabilities to the public, again without having informed St Jude Medical about the problems. In the event, the stock dipped marginally but not such that MedSec made significant profits from the venture.

Initially St. Jude denied the claims regarding vulnerabilities and argued that their software was secure. This appeared to be supported by researchers at the University of Michigan, who claimed to be unable to reproduce the same malfunctions found by MedSec. The same day, Muddy Waters released a video purportedly demonstrating some vulnerabilities, which may have been created using some bad assumptions about how the device should be configured or used [54]. St Jude Medical responded by bringing a law suit against MedSec in September 2016 [55, 56].

Independent research by Bishop Fox, published in October 2016, supported the claims of MedSec, agreeing that there were some vulnerabilities in the St. Jude systems [57]. In August 2017 the US Federal Drug Administration subsequently recalled 465,000 pacemakers manufactured by Abbott Laboratories, which had acquired St Jude Medical in January that year [54].

MedSec were criticised for working with Muddy Waters for occluding the central issue of their case. The independence of MedSec's research was brought into question through the possibility of their receiving financial reward for their findings. MedSec CEO Justine Bone responded that the company had deliberated over which course to take and concluded that the collaboration with Muddy Waters was the best option to force St Jude Medical into taking action. She claimed that St Jude had a poor history of responding to security flaws and referenced a reported case in which the company took two years to respond to a security flaw after learning of its existence. This history, she concluded, led MedSec to the conclusion that, "a partnership with Muddy Waters was the fastest route to improved product safety, improving patient safety and a better understanding of the risks faced by patients" [56].

On the one hand it seems as if MedSec were pre-judging St Jude Medical's likely response to the revelation of the security flaws. The grounds that St Jude Medical would not respond in a timely fashion appear weak and based on generalised industry behaviour and a rumour of foot-dragging in response to prior revelations. Furthermore, MedSec's motivations were brought into question by their decision to work with Muddy Waters to profit from shorting the stock. As David Robinson and Alex Halderman note, "researchers must be vigilant to retain as much independence as is feasible and transparent about the extent to which their end product is informed or shaped by other actors" [58]; p. 122).

On the other hand, MedSec's concerns regarding industry foot-dragging were not entirely misplaced. The traditional course of events, for independent security researchers to by-pass customers and inform vendors of flaws in their systems, has led to delays in patches being

developed and legal cases brought under the Digital Millennium Copyright Act and the Computer Fraud and Abuse Act against researchers. Furthermore, the independent research by Bishop Fox confirmed MedSec's claims, which were flatly denied by St Jude Medical [56]. Finally, VP of research at Veracode, Chris Eng has suggested that the MedSec precedent "has a lot of potential to be a net positive. We've all seen how consumer products are often designed and built in insecure ways, and let's face it, there has been virtually no improvement unless there's a major financial or reputational impact in doing so" [54].

Our interest in the MedSec/Muddy Waters/St Jude Medical case here is two-fold. In the first instance, should MedSec have partnered with an investment company to short the stock of a company it knew would suffer on the market once the flaws were made known? Secondly, should cybersecurity researchers be protected from legal action such as the attempt to sue MedSec pursued by St Jude? In both cases, the answers are unclear. In an ideal world perhaps, uncalled-for penetration testing would be "pure" of financial motives, but we do not live in such a world and pen testers, especially those who appear motivated to work in the public interest, need to be recompensed somehow. Likewise, researchers with genuine, public-spirited motivations should be protected from predatory practices by companies seeking to paper over cracks in their own security through legal action. That threat of legal action can be intimidating and serve as a chilling effect on legitimate research. While there are some developments in protecting researchers from prosecution, such as the CVD policy in The Netherlands [59], or DMCA exceptions in the US [60], this is not yet globally accepted practice. At the same time, it is entirely legitimate for companies to seek legal protection of IP and products. The concern here is not to exonerate cybersecurity researchers from legitimate legal action, but rather at the possibility of legal action being threatened to deter legitimate research which is in the public interest.

As is the nature of ethical dilemmas, there are no easy solutions to these issues. We suggest that the best way forward is the development of a code of conduct for cybersecurity research which establishes what is and is not acceptable behaviour under these circumstances. Such a code will not only serve to guide cybersecurity researchers who do not have the privilege of a university umbrella shielding them from legal action, but it will also provide support to those researchers who act within the guidelines and can thereby reasonably claim to be operating within the recognized ethical boundaries of the field of practice. We highlight below several ethical issues raised in security practice before turning to the issue of codes of conduct.

3.2. Summary of issues raised in cybersecurity practice

As noted in the Case Study 2, outside the university environment further ethical problems arise for those engaged in cybersecurity practice. These include many of the same issues faced by university research but are frequently complicated by a lack of institutional tradition and policy governing behaviour, an absence of an REB, and conflicts of interest between making money and doing "the right thing". The fact remains that a university exists for the development of knowledge aimed at furthering the public good. While many companies might see themselves fulfilling a similar role, for others the concern is not *endangering* the public good rather than seeking to further it.

As above, we break down the ethical issues related to cybersecurity practice into four broad areas in line with the Menlo principles: respect for persons, beneficence, justice, respect for law and public interest. Once more, this list is intended to be indicative rather than exhaustive.

3.2.1. Respect for persons

3.2.1.1. Informed consent. Informed consent remains a key issue for cybersecurity practice as it is for university research. However, practitioners often lack a tradition or relevant policies regarding the sourcing of informed consent, and in many cases this may be effectively impossible to obtain if, for instance, the user base exceeds a few thousand individuals. Furthermore, as noted in the Facebook/Cornell University emotional response research of January 2014, there may be a tradition among some groups, such as market research, of not seeking informed consent [61]. This then makes it difficult for the conscientious cybersecurity practitioner operating in these fields to insist on obtaining consent.

3.2.1.2. Trust. Trust is another area of concern, connecting the cybersecurity practitioner to those he or she is purportedly securing. There is an increasing recognition that security is best practiced through relationship with those secured rather than imposed upon them. An antipathetic relationship here is in no-one's interests, and yet security is often resented by employees while security teams often feel underappreciated [62]. Responses to this might involve increased transparency and access to cybersecurity teams, a focus on developing diversity within those teams, and efforts made by those teams in engaging with the workforce.

3.2.2. Beneficence

3.2.2.1. Privacy and control of data. As with university research, privacy and control of data are key issues in cybersecurity. Practitioners are likely to encounter personal data on a regular basis, whether they are interested in this or not, and they could be of a sensitive nature, such as data pertaining to bank or health records. The maintenance of privacy is thus crucial, and professional standards of confidentiality must be maintained. However, privacy is not the only issue at stake with personal data. There is also a central concern regarding the control of those data. For example, the recent scandal concerning Facebook and Cambridge Analytica was not primarily a matter of privacy, but of what was done with people's data, and the sharing of those data without consent [63]. Harms then emerge which go beyond the revelation of otherwise private information to the potential misuse of data (e.g. to influence election results) and may, if the data are handled ineptly, have a detrimental effect on the quality, integrity and future usability of those data.

A related concern is access to personal and/or sensitive data which may come accidentally through researching potential vulnerabilities in a related system. For example, if two networks are connected, then exploring a vulnerability in one may lead the researcher into the connected network (potentially unowned by the company employing the researcher) and through that to personal data [58]; p. 124). This would involve a clear privacy infraction, although not necessarily a violation on the part of the researcher if no intrusion into private data was intended.

Risk

A further ethical issue involves risk, and particularly questions of who is deciding on, as opposed to who is effected by, risky decisions, what are acceptable risk thresholds, and how risk is calculated [64]. As Wolff has demonstrated, there may be different ethical issues at stake as these vary between the decision-maker also being the cost-payer in a risky situation versus where the cost-payer is a person other than the decision-maker [65]. Furthermore, empirical research suggests that white males tend to tolerate higher levels of risk than women and non-white males, and that experts are more risk tolerant than the general public [66]. These suggest that current levels of risk acceptance may not be representative of society, and argue for a greater level of diversity in the decision-making process and for greater levels of public engagement [58]; pp. 120–28).

Security

There are obvious security-related issues at the heart of cybersecurity practice and if security is seen as an ethical issue, the maintenance of adequate security is itself an ethical issue. As such, compromises of security through insufficient funding, poor oversight of systems, late or no installation of patches, how and where data is stored, how that data is accessed, and poor training of staff in security awareness are all ethical concerns. Many of these issues may amount to professional negligence, such as the microsite server discovered at Greenwich University containing databases of nearly 20,000 people, including staff and students, which contained references to sensitive issues such as mental health which had not been updated or apparently even managed for 12 years [67,68].

At the same time, security is everyone's concern and so responsibility for security should not be seen to rest solely on the heads of those charged with its oversight. The security of an organisation (including cybersecurity) is the concern of all employees of that organisation. Training is an essential aspect in encouraging employee buy-in of security methods and systems. However, even after training, an employee may open an attachment in a spear-phishing email. Furthermore, limits in security budgets demand priorities are made, and as such some areas (including general staff training) are likely to be under-resourced or may be left to simplistic online training courses which do not encourage employees to take security seriously.

It is clear that the risks of cyberattack frequently fail to be understood, which may be the cause of many of the above problems. A lack of funding, poor training of staff, and a failure to install patches can all be hampered by a failure to recognize the immediacy or the gravity of the threat faced. One seemingly-obvious solution to this is to increase resources devoted to training, but this is not a panacea.

An alternative solution is to raise the profile of cybersecurity professionals within an organisation such that they are seen as a benefit rather than a burden. Hence if a marketing department decides to launch a new website to support an advertising campaign they could either rely on their own experience and training in cybersecurity or involve the cybersecurity team from the beginning of the project.

3.2.3. Justice

Bias

Diversity and related justice issues also extend beyond the gender and ethnic composition of cybersecurity teams to the impacts that cybersecurity efforts may have. For example, profiling behaviours outside cybersecurity practice has been demonstrated to embed bias in algorithmic code, and so similar attempts at profiling for the purposes of cybersecurity risk embedding similar discriminatory patterns [37,39]. A diverse composition of cybersecurity team members may help in the early identification of such patterns, as it might also in lowering risk thresholds.

Responsibility

There is an ongoing problem with cybersecurity insofar as the locus of responsibility is concerned [69]. This is less of a problem in academic research where the work is carried out under the auspices of an institution with its own REBs and structural hierarchies. However, in commercial research the locus of responsibility, and how far that responsibility extends, is unclear. Should a company be entirely responsible for developing its own cybersecurity? Is this so even when the company is (likely to be) subject to attack from foreign states or state-backed hackers? To what degree should the state take responsibility for protecting its own economy on the internet as it does in physical space, by providing safe places to trade?

3.2.4. Respect for law and public interest

3.2.4.1. Vulnerability disclosure. As demonstrated in Case Study 2, vulnerability disclosure is an issue for commercial as well as university-based cybersecurity research. When breaches occur, should these be reported, and to whom? On the one hand, sharing information increases vulnerability as one's defences become known, and one's experience of attacks shared. Yet on the other hand, it is arguably only by pooling experience that an effective defence can be mounted [69]; pp. 89–111). While this is undeniably risky, similar decisions made to share physical vulnerabilities in the past have led to positive developments in inter-corporate relationships and safety [70].

As noted in the MedSec case, there are conventions regarding the disclosure of discovered vulnerabilities and yet at least some corporations take advantage of these conventions. If there is a standard delay between disclosing a vulnerability to the company and disclosing it to the public, then the company may drag its heels in finding a solution. Furthermore, disclosure of the vulnerability to the public is damaging to the company or other companies using the same software (if a fix has not been found) and may increase awareness of the vulnerability, aiding future attacks. Robinson and Halderman note that there is a tension between those running vulnerable systems, for whom the *appearance* of a problem may be the greatest concern in terms of public relations, and those using the systems, for whom the problem *itself* is of far greater concern [58]; pp. 122–26).

A related problem is whether security researchers should agree to so-called gagging clauses which prevent them from publicly disclosing vulnerabilities which the company in question then refuses to address. Again, Robinson and Halderman argue that "researchers need to ensure that rules allow them to maintain their independence. If researchers are asked to sign a nondisclosure agreement, they should ensure that the terms allow them to disclose problems they might find, and do not overly restrict their ability to perform future work" [58]; p. 123). In the same chapter, the authors also consider challenges regarding the timing of disclosure of vulnerabilities in e-voting machines: too early and the researchers risk disrupting an election; too late and they risk damaging trust in the election result [58]; p. 126).

Finally, the manner in which the company engaging in security research or audits responds to revelations either of flaws or a lack of discovered vulnerabilities is a concern. We have already considered companies which do not attempt to address flaws and may impose gagging clauses on researchers (or threaten legal action). However, there are also issues in companies taking negative results as demonstrations that their products are safe. It is clearly not the case that the failure of one researcher (or even many researchers) to find a flaw implies that there are no such flaws. Security researchers should be aware of how their findings will be interpreted and used.

Business ethics

Business ethics, and the associated conflicts that arise specifically because of competing interests in security and making money, do not fit easily within the Menlo framework. Security should not be ignored in the interests of channelling funds into profit-making activities, and a minor degree of prescience will suggest that good security will strengthen a company's reputation and client trust in that organisation. Nonetheless, it would be naive to suggest that conflicts of interest do not emerge between individual interests, public interests and corporate interests. A matter of days before the 2017 Equifax breach was made public, certain senior executives sold their shares in the company. Two subsequently pleaded guilty to insider trading while others denied any wrongdoing [71]. In Case Study 2, MedSoft were not acting on information that was only available to those inside the company and so the decision to short the stock was not a matter of insider trading. Nonetheless, the wisdom of such a move might be questioned, as might the

decision to publicize the vulnerabilities rather than approach St Jude first, giving them time to develop patches. Finally, the decision of Marissa Meier, then CEO of Yahoo, not to inform the public of the hacks in 2013 and 2014 regarding 3bn accounts seems hard to understand in terms other than the attempted protection of the company's image at the expense of users' security [72,73]. As such, concerns of business ethics should form a central plank in ethical assessment of cybersecurity research in the practitioner community.

4. Existing guidelines and recommendations

The aforementioned ethical issues are legion and complicated, albeit hardly new to the cybersecurity community. Despite this, there is relatively little guidance as to how practitioners should proceed in many of these cases. The Association of Internet Researchers has produced guidelines for ethical research, which are currently on their second edition [74]. However, while valuable in themselves, these present a series of questions for reflection rather than principles for guiding conduct. A similar approach is taken by Networked Systems Ethics, a project from the Oxford Internet Institute [75] in that questions are raised to initiate a process rather than provide clear guidelines. While this approach taken in both frameworks is laudable it is arguably of limited value to those seeking an understanding of what they should or should not be doing in researching cybersecurity. Furthermore, both are targeted at the broader ethical concerns of internet research, rather than cybersecurity research *per se*. As noted above, there are several issues which arise in cybersecurity research (such as vulnerability disclosure) which are not generally recognized as issues in general internet research. The challenge, as recognized by the SIGCOMM 22015 Program Committee is that, "The controversy [surrounding Encore] arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of [at least some] experiments" [1].

An alternative approach is the development of codes of conduct, either at an institutional level, such as through the IEEE and ACM, or at the corporate level [76,77]. While the IEEE and ACM have each developed codes of conduct, these are again not designed with cybersecurity primarily in mind. The Menlo principles, which are focused on cybersecurity research, are a helpful start but very broad. While they can address many if not all the issues raised in this paper, the guidance provided in some cases, such as vulnerability disclosure, is not always clear. Furthermore, codes of conduct must be supported by effective sanctions if they are broken. Codes can be incredibly powerful tools in the hands of professionals seeking to resist pressure to act unethically [78]. However, for them to function effectively, Davis notes that all professionals need to adhere to the code in the face of adversity. To encourage this, the professional bodies presenting such codes must ensure that they are supported, and transgressions punished. It is this which makes a code of conduct so powerful in medicine (where physicians can be struck off for unethical behaviour) and so weak in other professions where the code may be routinely ignored by a significant minority of practitioners. Such a support system does not exist in the field of cybersecurity, and we feel that the field is worse off for it.

Such guidelines will be of use for both practitioners and those in academic research. In addition to this, academic REBs need to reflect on their constitution and whether they are sufficiently competent at present to manage the issues arising in this paper. Ideally, a generation of computer scientists trained in ethics at the undergraduate and postgraduate level would be members of any such committee. When these are not available, though, REBs should ensure that they are able to draw on the experience of computer scientists for decisions. Furthermore, the aid provided by an REB is often a significant benefit for researchers looking to practice ethically. Professional bodies such as the ACM and IEEE could also look to provide research ethics recommendations for members.

Finally, there is a clear need for the development of an active conversation regarding ethics in the research and practice of cybersecurity. This, too, is lacking, owing in part to the relative paucity of ethics teaching provided to computer scientists in higher education, especially when it comes to teaching the ethics of cybersecurity [52]. While there are attempts to address this [79–81], these are recent and, as the case studies demonstrate, need to gain wide traction rapidly. It is notable that UK degree courses accredited by the British Computer Society as of January 2020 have an obligation to "give students an awareness of external factors which may affect the work of the computer professional. These may vary according to the orientation of the programme and the likely destination of students, but examples could include … computer security" [82]. Furthermore, as Hughes et al. note, these concerns are not limited to academic education in the global North but are prevalent throughout computer science teaching worldwide [83]. Through the publication of this paper we hope to stimulate further discussion at the academic and practitioner level regarding the ethical issues raised here, and doubtless others that we have not addressed here.

5. Conclusion

In this paper we have argued that current methods of oversight and guidance regarding cybersecurity ethics are inadequate. We have considered these methods in two areas: university-based development and the community of practising experts.

In the former we argued that the problems stem from a lack of awareness among members of ethical review committees as to the nature of relevant ethical problems, such as considered in Case Study 1. In the latter there is a lack of adequate guidance or accountability which forms a barrier to consistent ethical practice, illustrated in Case Study 2. We have therefore argued that there needs to be a greater appreciation of the risks of cybersecurity development in academic ethical review committees and clear (and enforceable) codes of conduct for, or at least active discourse within, the professional community which cover development and practice.

CRediT authorship contribution statement

Kevin Macnish: Conceptualization, Formal analysis, Methodology, Investigation, Project administration, Writing - original draft, Writing - review & editing. **Jeroen van der Ham:** Conceptualization, Formal analysis, Methodology, Investigation, Project administration, Writing - original draft, Writing - review & editing.

References

[1] S. Burnett, N. Feamster, Encore: lightweight measurement of Web censorship with cross-origin requests, in: Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM '15. ACM, New York, NY, USA, 2015, pp. 653–667, https://doi.org/10.1145/2785956.2787485.
[2] J.W. Byers, Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests – Public Review, Technical Report, Department of Computer Science, Boston University, 2015, http://conferences.sigcomm.org/sigcomm/2015/pdf/reviews.
[3] T.L. Beauchamp, J.F. Childress, Principles of Biomedical Ethics, sixth ed., OUP USA, New York, 2009.
[4] Economic, Social Research Council, Our core principles - economic and social research council [WWW Document]. URLaccessed 1.11.19, https://esrc.ukri.org/funding/guidance-for-applicants/research-ethics/our-core-principles/, 2019.
[5] R.R. Faden, T.L. Beauchamp, The concept of informed consent, in: P. of P. Beauchamp, L. Walters, J.P. Kahn, A.C. Mastroianni (Eds.), Contemporary Issues in Bioethics, Cengage Learning, Boston, MA, 2008, pp. 166–170.
[6] R.J. Levine, Informed consent: some challenges to the universal validity of the Western model, in: T. Beauchamp, L. Walters, J.P. Kahn, A.C. Mastroianni (Eds.), Contemporary Issues in Bioethics, Cengage Learning, Boston, MA, 2008, pp. 170–175.
[7] K. Macnish, Informed consent, in: C. Veliz (Ed.), Data, Privacy and the Individual, IE University Press, Madrid, 2019.
[8] N. Manson, O. O'Neill, Rethinking Informed Consent in Bioethics, 2007 (Cambridge).
[9] University of Oxford, Policy on the ethical conduct of research involving human participants and personal data, Research Support [WWW Document]. URLaccessed

1.11.19, https://researchsupport.admin.ox.ac.uk/governance/ethics/committees/policy#collapse395071, 2015.
[10] World Medical Association, G.A. of the W.M, World Medical Association Declaration of Helsinki: ethical principles for medical research involving human subjects, J. Am. Coll. Dent. 81 (2014) 14.
[11] J. Aycock, E. Buchanan, S. Dexter, D. Dittrich, Human subjects, agents, or bots: current issues in ethics and computer security research, in: G. Danezis, S. Dietrich, K. Sako (Eds.), Financial Cryptography and Data Security, Springer, Berlin, 2012, pp. 138–145.
[12] M.L. Johnson, S.M. Bellovin, A.D. Kromytis, Computer security research with Huma subjects: risks, benefits and informed consent, in: G. Danezis, S. Dietrich, K. Sako (Eds.), Financial Cryptography and Data Security, Springer, Berlin, 2012, pp. 131–137.
[13] J. van der Ham, Embedding ethics in system administration education, The USENIX Journal of Education in System Administration 1 (2015) 1–9.
[14] D. Dittrich, E. Kenneally, The Menlo report: ethical principles guiding information and communication technology research [WWW Document]. CAIDA. URLaccessed 3.11.19, http://www.caida.org/publications/papers/2012/menlo_report_actual_formatted/index.xml, 2012.
[15] E. Department of Health, The Belmont Report. Ethical principles and guidelines for the protection of human subjects of research, J. Am. Coll. Dent. 81 (2014) 4.
[16] J.M. Sims, A brief review of the Belmont report, Dimens. Crit. Care Nurs. 29 (2010) 173–174.
[17] G.J. Annas, M.A. Grodin, Th Nuremberg code, in: E.J. Emanuel, C.C. Grady, R. A. Crouch, R.K. Lie, F.G. Miller, D.D. Wendler (Eds.), The Oxford Textbook of Clinical Research Ethics, Oxford University Press, 2008, pp. 136–140.
[18] N. Code, The Nuremberg code, Trials of war criminals before the Nuremberg military tribunals under Control Council Law 10 (1949) 181–182.
[19] A. Dhai, The research ethics evolution: from Nuremberg to Helsinki, S. Afr. Med. J. 104 (2014) 178–180.
[20] E. Shuster, Fifty years later: the significance of the Nuremberg Code, N. Engl. J. Med. 337 (1997) 1436–1440.
[21] T.L. Beauchamp, Autonomy and consent, in: F. Miller, A. Wertheimer (Eds.), The Ethics of Consent: Theory and Practice, OUP USA, Oxford, New York, 2009, pp. 55–78.
[22] T. Moore, R. Clayton, Ethical dilemmas in take-down research, in: G. Danezis, S. Dietrich, K. Sako (Eds.), Financial Cryptography and Data Security, Springer, Berlin, 2012, pp. 146–153.
[23] K. Macnish, Privacy in research ethics, in: R. Iphofen (Ed.), Handbook of Research Ethics and Scientific Integrity, Springer International Publishing, 2020.
[24] A.S. Elmaghraby, M.M. Losavio, Cyber security challenges in Smart Cities: safety, security and privacy, J. Adv. Res. 5 (2014) 491–497.
[25] N. Kshetri, Blockchain's roles in strengthening cybersecurity and protecting privacy, Telecommun. Pol. 41 (2017) 1027–1038.
[26] C. Landwehr, D. Boneh, J.C. Mitchell, S.M. Bellovin, S. Landau, M.E. Lesk, Privacy and cybersecurity: the next 100 years, Proc. IEEE 100 (2012) 1659–1673.
[27] J. Liu, Y. Xiao, S. Li, W. Liang, C.P. Chen, Cyber security and privacy issues in smart grids, IEEE Communications Surveys & Tutorials 14 (2012) 981–997.
[28] B.J. Murrill, E.C. Liu, R.M. Thompson, Smart Meter Data: Privacy and Cybersecurity, Congressional Research Service, 2012 (Library of Congress).
[29] EU Parliament, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), 2016 (OJ L).
[30] D.A. Brown, A.N. Hasso, Toward a uniform policy for handling incidental findings in neuroimaging research, Am. J. Neuroradiol. 29 (2008) 1425–1427, https://doi.org/10.3174/ajnr.A1227.
[31] W. Burke, A.H. Matheny Antommaria, R. Bennett, J. Botkin, E.W. Clayton, G. E. Henderson, I.A. Holm, G.P. Jarvik, M.J. Khoury, B.M. Knoppers, N.A. Press, L. F. Ross, M.A. Rothstein, H. Saal, W.R. Uhlmann, B. Wilfond, S.M. Wolf, R. Zimmern, Recommendations for returning genomic incidental findings? We need to talk! Genetics in Medicine 15 (2013) 854–859, https://doi.org/10.1038/gim.2013.113.
[32] S.S. Kalia, K. Adelman, S.J. Bale, W.K. Chung, C. Eng, J.P. Evans, G.E. Herman, S. B. Hufnagel, T.E. Klein, B.R. Korf, K.D. McKelvey, K.E. Ormond, C.S. Richards, C. N. Vlangos, M. Watson, C.L. Martin, D.T. Miller, Recommendations for reporting of secondary findings in clinical exome and genome sequencing, 2016 update (ACMG SF v2.0): a policy statement of the American College of Medical Genetics and Genomics, Genet. Med. 19 (2017) 249–255, https://doi.org/10.1038/gim.2016.190.
[33] A.L. McGuire, S. Joffe, B.A. Koenig, B.B. Biesecker, L.B. McCullough, J. S. Blumenthal-Barby, T. Caulfield, S.F. Terry, R.C. Green, Ethics and genomic incidental findings, Science 340 (2013) 1047–1048, https://doi.org/10.1126/science.1240156.
[34] J. Viberg, M.G. Hansson, S. Langenskiöld, P. Segerdahl, Incidental findings: the time is not yet ripe for a policy for biobanks, Eur. J. Hum. Genet. 22 (2014) 437–441, https://doi.org/10.1038/ejhg.2013.217.
[35] S.M. Wolf, G.J. Annas, S. Elias, Patient Autonomy and incidental findings in clinical genomics, Science 340 (2013) 1049–1050, https://doi.org/10.1126/science.1239119.
[36] D.B. Resnik, P.R. Finn, Ethics and phishing experiments, Sci. Eng. Ethics 24 (2018) 1241–1252, https://doi.org/10.1007/s11948-017-9952-9.
[37] K. Macnish, Unblinking eyes: the ethics of automating surveillance, Ethics Inf. Technol. 14 (2012) 151–167, https://doi.org/10.1007/s10676-012-9291-0.
[38] B.D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, L. Floridi, The ethics of algorithms: mapping the debate, Big Data & Society 3 (2016), https://doi.org/10.1177/2053951716679679, 2053951716679679.
[39] C. O'Neil, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy, Crown/Archetype, 2016.
[40] A. Arora, R. Telang, H. Xu, Optimal policy for software vulnerability disclosure, Manag. Sci. 54 (2008) 642–656.
[41] J.P. Choi, C. Fershtman, Internet security, vulnerability disclosure and software provision, 2005.
[42] A. Hahn, M. Govindarasu, Cyber Vulnerability Disclosure Policies for the Smart Grid, 2012 IEEE Power and Energy Society General Meeting, 2012, pp. 1–5. IEEE.
[43] A. Kuehn, M. Mueller, Shifts in the cybersecurity paradigm: zero-day exploits, discourse, and emerging institutions, in: Proceedings of the 2014 New Security Paradigms Workshop, ACM, 2014, pp. 63–68.
[44] B.C. Panton, Strengthening US DoD cyber security with the vulnerability market, 2013 (AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH GRADUATE SCHOOL OF …).
[45] P.P. Swire, A model for when disclosure helps security: what is different about computer and network security, J. Telecommun. High Technol. Law 3 (2004) 163.
[46] C. Zheng, Y. Zhang, Y. Sun, Q. Liu, IVDA: international vulnerability database alliance, in: 2011 Second Worldwide Cybersecurity Summit (WCS), IEEE, 2011, pp. 1–6.
[47] C. Soghoian, Enforced community standards for research on users of the tor anonymity network, in: G. Danezis, S. Dietrich, K. Sako (Eds.), Financial Cryptography and Data Security, Springer, Berlin, 2012, pp. 146–153.
[48] M. Allman, What ought a program committee to do? WOWCS 8 (2008) 1–5.
[49] S.L. Garfinkel, IRBs and security research: myths, facts and mission creep, in: UPSEC'08 Proceedings of the 1st Conference on Usability, Psychology, and Security, USENIX Association, Berkeley, CA, 2008.
[50] E.A. Buchanan, C.M. Ess, Internet research ethics and the institutional review board: current practices and issues, SIGCAS Comput. Soc. 39 (2009) 43–49, https://doi.org/10.1145/1713066.1713069.
[51] E. Buchanan, J. Aycock, S. Dexter, D. Dittrich, E. Hvizdak, Computer science security research and human subjects: emerging considerations for research ethics boards, Journal of Empirical Research on Human Research Ethics 6 (2011) 71–83, https://doi.org/10.1525/jer.2011.6.2.71.
[52] C. Fiesler, N. Garrett, N. Beard, What do We teach when We teach tech ethics? A syllabi analysis, in: Proceedings of the 51st ACM Technical Symposium on Computer Science Education, SIGCSE '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 289–295, https://doi.org/10.1145/3328778.3366825.
[53] S. Baase, T.M. Henry, A Gift of Fire: Social, Legal, and Ethical Issues for Computing Technology, 5 edition, Pearson, NY, NY, 2017.
[54] T. Spring, Researchers: MedSec, Muddy Waters set bad precedent with St. Jude medical short, The first stop for security news | Threatpost. URLaccessed 7.4.18, https://threatpost.com/researchers-medsec-muddy-waters-set-bad-precedent-with-st-jude-medical-short/120266/, 2016.
[55] S. Nichols, St Jude sues short-selling MedSec over pacemaker "hack" report [WWW Document]. The Register. URLaccessed 7.4.18, https://www.theregister.co.uk/2016/09/07/st_jude_sues_over_hacking_claim/, 2016.
[56] J.C. Weigelt, St. Jude medical brings legal action against Muddy Waters and MedSec [WWW Document]. URLaccessed 7.4.18, http://media.sjm.com/newsroom/news-releases/news-releases-details/2016/St-Jude-Medical-Brings-Legal-Action-Against-Muddy-Waters-and-MedSec/default.aspx, 2016.
[57] C.D. Livitt, Preliminary expert report of carl D. Livitt (No. 0:16-cv-03002-DWF-JSM), 2016.
[58] D.G. Robinson, J.A. Halderman, Ethical issues in E-voting security analysis, in: G. Danezis, S. Dietrich, K. Sako (Eds.), Financial Cryptography and Data Security, Lecture Notes in Computer Science, Springer, Berlin, 2012.
[59] National Cyber Security Centre, Coordinated vulnerability disclosure: the guideline | NCSC (webpagina), 2018.
[60] A. Alva, DMCA security research exemption for consumer devices [WWW Document]. Federal Trade Commission. URLaccessed 7.1.19, https://www.ftc.gov/news-events/blogs/techftc/2016/10/dmca-security-research-exemption-consumer-devices, 2016.
[61] R. Meyer, Everything We know about facebook's secret, Mood Manipulation Experiment [WWW Document]. The Atlantic. URLaccessed 7.10.18, https://www.theatlantic.com/technology/archive/2014/06/everything-we-know-about-facebooks-secret-mood-manipulation-experiment/373648/, 2014.
[62] E. Tucker, Cyber security – why you're doing it all wrong [WWW Document]. ComputerWeekly.com. URLaccessed 12.17.18, https://www.computerweekly.com/opinion/Cyber-security-why-youre-doing-it-all-wrong, 2018.
[63] M. Ienca, E. Vayena, Cambridge Analytica and online manipulation [WWW document]. Scientific American blog network, URLaccessed 7.10.18, https://blogs.scientificamerican.com/observations/cambridge-analytica-and-online-manipulation/, 2018.
[64] S.O. Hansson, The Ethics of Risk: Ethical Analysis in an Uncertain World, Palgrave Macmillan, 2013.
[65] J. Wolff, Five types of risky situation, Law, Innovation and Technology 2 (2010) 151–163, https://doi.org/10.5235/175799610794046177.
[66] H. Hermansson, Towards a fair procedure for risk management, J. Risk Res. 13 (2010) 501–515, https://doi.org/10.1080/13669870903305903.
[67] J.E. Dunn, Server? What server? Site forgotten for 12 years attracts hacks, fines URL Naked Security (2018), 2018. https://nakedsecurity.sophos.com/2018/05/22/server-what-server-site-forgotten-for-12-years-attracts-hacks-fines/. (Accessed 3 November 2019).

[68] Information Commissioner's Office, The University of Greenwich fined £120,000 by Information Commissioner for "serious" security breach [WWW Document]. URL3.11.19, https://icoumbraco.azurewebsites.net/about-the-ico/news-and-events/news-and-blogs/2018/05/the-university-of-greenwich-fined-120-000-by-information-commissioner-for-serious-security-breach/, 2018.
[69] A.N. Guiora, Cybersecurity: Geopolitics, Law, and Policy, 1 edition, Routledge, Boca Raton, FL, 2017.
[70] J. Turnbull, Ethics and employability, in: R. Lawlor (Ed.), Engineering in Society: beyond the Technical, University of Leeds, Leeds, 2015.
[71] K. Brumback, A former Equifax manager just got sentenced for insider trading after making $75,000 off of its massive data breach [WWW Document]. Business Insider. URLaccessed 3.11.19, https://www.businessinsider.com/ex-equifax-exec-sentenced-for-insider-trading-after-massive-data-breach-2018-10, 2018.
[72] N. Stone, The Yahoo Cyber Attack & what should you learn from it? [WWW Document]. Cashfloat. URLaccessed 12.17.18, https://www.cashfloat.co.uk/blog/technology-innovation/yahoo-cyber-attack/, 2017.
[73] S. Thielman, Yahoo Hack: 1bn Accounts Compromised by Biggest Data Breach in History, The Guardian, 2016.
[74] A. Markham, E. Buchanan, Ethical decision-making and internet research: version 2.0. recommendations from the AoIR ethics working committee, Available online: aoir.org/reports/ethics2.pdf, 2012.
[75] Oxford Internet Institute, Networked systems ethics [WWW Document]. URLaccessed 3.11.19, http://networkedsystemsethics.net/index.php?title=Networked_Systems_Ethics_-_Guidelines, 2017.
[76] ACM, ACM code of ethics and professional conduct [WWW Document]. URLaccessed 3.11.19, https://www.acm.org/code-of-ethics, 2018.
[77] IEEE, IEEE code of ethics [WWW Document]. URLaccessed 3.11.19, https://www.ieee.org/about/corporate/governance/p7-8.html, 2019.
[78] M. Davis, Thinking like an engineer: the place of a code of ethics in the practice of a profession, Philos. Publ. Aff. (1991) 150–167.
[79] G. Austin, Cyber Security Education: Principles and Policies, Routledge, 2020.
[80] Joint Task Force (JTF) on Cybersecurity Education, Cybersecurity curricula 2017, curriculum guidelines for post-secondary degree programs in cybersecurity, a report in the computing curricula series, ACM/IEEE-CS/AIS SIGSEC/IFIP WG 11 (2017) 8. Version 1.0, New York.
[81] D. Shoemaker, A. Kohnke, K. Sigler, The Cybersecurity Body of Knowledge: the ACM/IEEE/AIS/IFIP Recommendations for a Complete Curriculum in Cybersecurity, CRC Press, 2020.
[82] BCS, Guidelines on Course Accreditation, British Computer Society London, 2020.
[83] J. Hughes, E. Plaut, F. Wang, E. von Briesen, C. Brown, G. Cross, V. Kumar, P. Myers, Global and local agendas of computing ethics education, in: Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education, ITiCSE '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 239–245, https://doi.org/10.1145/3341525.3387423.
[84] K.E. Himma, The ethics of tracing hacker attacks through the machines of innocent persons, International Journal of Information Ethics 2 (2004) 1–13.
[85] D. Ross, The Right and the Good, New edition, Clarendon Press, 2002.
[86] J. Rawls, A Theory of Justice, Revised edition, Harvard University Press, 1999.
[87] A. Narayanan, B. Zevenbergen, No Encore for Encore? Ethical Questions for Web-Based Censorship Measurement (SSRN Scholarly Paper No. ID 2665148), Social Science Research Network, Rochester, NY, 2015.