

Digitalisierung und Verantwortung: Ambivalenz von Öffentlichkeit am Beispiel von Forschung

Dual-Use in Cybersecurity Research. Towards a New Culture of Research Ethics

1 Introduction

The fact that information and communication technologies (ICTs) increasingly shape our online and offline lifeworlds has led to the emergence of a new societal threat in the form of vulnerabilities in critical ICT systems that may be exploited by malicious actors.¹ Cybersecurity researchers work on finding such vulnerabilities and on identifying new attack vectors, i.e. they systematically step into the role of attackers. Normatively, however, the goal of this research is to strengthen ICTs against cyberattacks and, thus, to reduce the societal threat.

However, this implies a dual-use potential: as in any discipline, results in cybersecurity research need to be published at some point and the disclosure of vulnerabilities to software companies and ICT administrators is an integral part of ensuring that ICT vulnerabilities are closed; such a disclosure, however, can be misused by attackers, as well. Thus, instruments for ethical orientation in regard of these dual-use issues seem to be necessary to identify potential issues and to guide researchers in dealing with them. While corresponding resources, such as codes of conduct, already exists for cybersecurity research, we will show that they currently do not adequately address the need for ethical orientation in the academic context. This, we argue, creates the need for a new research ethics culture in cybersecurity research.

We will start our argument by identifying why the risk of cyberattacks is structurally different to traditional security threats, as they have a scaling risk dynamic (section 2). This is why we cannot rely on established forms of security production such as policing and, instead, involve cybersecurity research to answer this challenge. We will then show that this implies dual-use issues, however, as published findings may be misused (section 3). Finally, we will describe the steps that have so far been developed in the field, but we also argue that these

1 The research for this article was conducted as part of the Graduate Academy “SecHuman – Security for Humans in Cyberspace“, which is funded by the state of North Rhine-Westphalia, Germany. The article is a more concise and translated version of Weydner-Volkmann and Cassing (2023).

steps do not provide sufficient orientation to deal with situations typical for cybersecurity research (section 4). Hence, we conclude that cybersecurity needs to take further steps towards a professionalized culture of research ethics, an *Ethics of Cybersecurity* that may incorporate findings from applied ethics and technology assessment (TA).

2 The production of security through cybersecurity research

One example for the constant evolvement of new individual, economic and social vulnerabilities due to digitalization are ransomware attacks. Here, after exploiting an ICT vulnerability, attackers encrypt data and extort a ransom in exchange for handing over the key that is needed to restore the data. The U.S. Federal Bureau of Investigation (FBI) notes that the number of ransomware attacks decreased in 2022, but that "ransomware remains a serious threat to the public and to [...] economy" (FBI 2022, p. 3). Out of all forms of cybercrime, the German Federal Police (*Bundeskriminalamt*, BKA) considers ransomware attacks to currently have greatest potential for societal damage (cf. BKA 2020, p. 22). Ransomware attacks may affect individuals, but it can just as easily affect critical infrastructure, as in the case of the University Hospital of Düsseldorf in 2020 (cf. BKA 2020, p. 26; Silomon 2020). Hence, today, cybersecurity is a central societal concern – and the identification and closure of vulnerabilities within cybersecurity research plays a vital role (cf. Wagner 2020, p. 116).

The need for more extensive research efforts in TA on ICT vulnerabilities has already been highlighted (Weber et al. 2020). Still, there is a clear lack of publications dealing in more detail with the *ethical* implications of dealing with vulnerabilities in ICT systems and the role of cybersecurity research. In other words: With few notable exceptions (e.g., Christen et al. 2020; Dunn Cavely 2014; Macnish/van der Ham 2020), the ethical perspective, particularly with respect to the societal dimension of developing dual-use techniques and technologies, is largely absent in cybersecurity, but also in TA discussions.

One may wonder, however, if there really is a need for an ethical discourse specifically on issues in cybersecurity research. After all, one may argue that established ways to deal with other technological risks and their ensuing ethical implications can be applied seamlessly to the cybersecurity context. To address this, we will distinguish two types of established risk domains: (1) technical "safety" questions like the operational reliability, accident prevention, or failure handling, which have long been addressed in TA, and (2) "security" aspects

that deal with malicious actors, crime, terrorism and deliberate attacks. Here, the societal impact of introducing surveillance and control technologies for the production of security has been addressed in Surveillance Studies and Security Ethics. In a simplistic manner, one could say that the ethical issues for (1) mostly deal with the question of how safe is safe enough (e.g., with regard to the operation of nuclear power plants), while the ethical issues discussed for security technologies (2) have mostly been discussed in terms of power structures and impact on values such as privacy or non-discrimination.

What becomes clear in regard to ICT vulnerabilities is that they neither cleanly fit the safety category nor the security category. This is because typical safety requirements for ICTs encompass statistically quantifiable reliability parameters. These parameters indicate the system's ability to function without malfunctions for a specified period (cf. Eusgeld et al. 2008, p. 59). Ignoring deliberate attacks, most ICT vulnerabilities would not pose safety concerns. On the other hand, although cybersecurity research presupposes malicious actors, the developed technologies and techniques do not target attackers (like surveillance and control technologies), but rather the *robustness of ICTs* against deliberate attacks. Thus, while techniques and technologies that deal with finding vulnerabilities clearly fall within the realm of security issues rather than safety issues, established approaches of ethical and societal reflection (within or outside of TA) fail to properly address the nature of the dual-use problematic in cybersecurity research.

Still, one may wonder what prevents us from applying traditional policing approaches in the context of cybersecurity. We propose that ICT vulnerabilities in digitized societies introduce a changed risk dynamic as attacks benefit from “scaling effects” known from the economic context. Three dimensions can be distinguished:

1. *Spatially*, in the case of classic security problems, attacks are largely localized. One may consider a mundane bicycle theft – here, a thief needs to gain physical access to a bike. In contrast, exploiting a vulnerability in networked ICTs, e.g. for a ransomware attack, may be carried out from almost any point on the globe.
2. *Temporally*, finding a vulnerability and exploiting it for a ransomware attack may certainly take a lot of time – probably even longer than picking the lock of a bike. However, while every single bicycle theft now takes a similar amount of time, follow-up attacks on ICTs can often be fully automated and thus carried out *en masse* with ever decreasing effort.

3. *Topologically*, many forms of crime presuppose some form of a social relationship between the attacker and the victim – contrary to cybercrime. To continue the analogy between bicycle theft and ransomware attack, we have to extend the example: a stolen bicycle in this scenario is not being resold, but rather returned to the owner for a ransom. For this to work, the attacker needs some information about the owner, for example who they are or how much the bike is worth to the owner. Attacks on ICTs, on the other hand, are typically carried out in mutual anonymity and the relationship can be virtually random.

Due to the interaction of these three dimensions, attacks on ICT systems often develop a *scaling* dynamic: they can already be carried out globally, anonymously and *en masse* by individuals and small groups. Hence, policing cybercrime faces major hurdles in this regard. This is why early detection and (dis)closure of vulnerabilities have such a high societal value: the occurrence of cyberattacks even on critical infrastructure is considered a fact of modern life and so is the assumption that all complex ICTs have some vulnerabilities. Hence, the production of cybersecurity cannot focus on the attacker, but needs to focus on the robustness of the systems. Consequently, from a societal perspective, a different set of actors is given lead roles in the production of cybersecurity: researchers (cf. BMBF 2020; Wagner 2020, p. 116).

For those researchers, however, to explore the robustness of ICTs means that they themselves act like attackers, i.e. cybersecurity researchers fulfil their societal role by demonstrating the vulnerability of ICTs in a replicable manner, thereby indirectly contributing to safer systems. As argued above and as will be explored in more detail in the next section, this entails ethical challenges in the form of dual-use issues, especially with regard to the scientific publication of research results on vulnerabilities.

3 Disclosing research results as a dual-use issue

As mentioned earlier, a fundamental dilemma for researchers arises when they publish their research results: the publication of research results is an essential part of the academic work process and scientific progress, and as such it is covered by the concept of academic freedom in many countries (cf. Kováts/Rónay 2023, p. 4; Reydon 2013, p. 68). Sharing research results is essential in order to give the scientific community the opportunity to review and build on the results and, not least, publications are central for building a scientific career. At

the same time, however, the publication poses a societal risk as knowledge can be misused by third parties if it describes vulnerabilities, methods for detecting vulnerabilities, or the fact that a particular system is insecure.

Thus, researchers must weigh whether and to whom disclosing a result could lead to substantial insecurity for those dependent on the operation of an ICT. The weighing has to reflect that a disclosing publication of vulnerabilities may very well lead to containment of risks: the owner of the affected system (or other responsible actors) may react by developing and distributing patches or updates; users may act particularly cautious when they know about a certain vulnerability. It therefore remains somewhat ambivalent whether disclosing vulnerabilities ultimately increases societal risks or produces security. Currently, this reflective weighing is almost entirely in the hands of the researchers,² leaving them alone with the responsibility to assess the extent of the complex of ambiguous risks for individuals, society or companies.

Which criteria should researchers use to weigh these risks? Who should be responsible for which outcomes? There can be no general answer to these questions and, thus, there is a need for situational ethical orientation. In cybersecurity research, we can observe approaches that already try to address this need. Ethical codes of conduct and best practices, among others, have been developed to guide researchers in ethical matters. Despite these efforts, however, the need for situational ethical orientation remains, as we will see in the following section.

4 Current forms of orientation

As indicated above, it can be observed that the current research culture in cybersecurity hardly addresses the field's dual-use issues as part of a systematic theory-based research debate. As part of the research practice, however, at least four approaches were developed due to ethical conflicts. A first approach is the formulation of ethical codes of conduct. There are various codes in the broader field of IT (e.g., *ACM Code of Ethics and Professional Conduct* (ACM 2018) or *IFIP Code of Ethics and Professional Conduct* (IFIP 2020)). They have in

2 For the sake of completeness, it should be mentioned that the decision can be influenced by multiple factors. It can, for example, damage a company's reputation, reduce customer confidence, and increase the risk of misuse of the vulnerability if it becomes known that a company operates (potentially) vulnerable ICTs (cf. Dreißigacker et al. 2020, p. 150). Therefore, affected companies have a great interest to not disclose a vulnerability and they sometimes threaten legal action not to disclose vulnerabilities.

common that they formulate principles, have no juridical binding and depend on researchers' self-commitment. In the field of cybersecurity, the so-called Menlo Report (cf. Dittrich/Kenneally 2012) is of central importance. It is addressed to various professional groups that explicitly encompass researchers (ibid., p. 5). However, research ethical issues are not addressed in detail and the codes are hard to apply to real scenarios (cf. Macnish/van der Ham 2020, p. 8). The publication of vulnerabilities, for example, is discussed only by formulation of the principle that it needs to benefit society after deliberation of the pros and cons (cf. Dittrich/Kenneally 2012, p. 11). More concrete recommendations that could guide this deliberation are lacking. Here, one of the central issues of principlism surfaces: The application of the principles presupposes a certain ethical competence. At the same time, there is a lack of institutionalized structures in the field ensuring that this level of competence (or even a certain sensitivity to such dual-use issues) is imparted. Despite these practical challenges, the Menlo Report is the field's standard reference for ethical orientation.

As a second approach to research ethics, one could point towards the emergence of research ethics boards at major cybersecurity conferences, where the report's principles are often highlighted. Conferences, not journals, are the central publication medium in the field of cybersecurity research. Submissions to the conference are peer-reviewed before they are published. While the peer-review mainly concerns the *technical* quality of the submission, more and more attention is also being paid to ethical aspects (cf. Usenix 2021; NDSS 2022). For this purpose, ethical boards are being established and consulted whenever reviewers flag ethical concerns for submissions. One example of the formation of such a board is the Institute of Electrical and Electronics Engineering's (IEEE) cybersecurity conference, which is one of the field's most renowned. Here, the establishment of an ethics board and the formulation of a corresponding ethical code was initiated after a prominent submission had been deemed to violate the principle of informed consent and thereby ignited a discussion on standardizing research ethical requirements ("Hypocrite commits paper"; Loschwitz 2021; Salter 2021; Vaughan-Nichols 2021). For the committee and the conference organizers, the Menlo Report (cf. Dittrich/Kenneally 2012) was and is the main point of reference (IEEE 2022).

With regard to ethics, this raises similar concerns as mentioned above: to what extent can a principlism-based ethics code serve as an effective basis for orientation in the field? It stands to reason that extended ethics catalogues will not provide enough orientation for researchers that face typical dual-use issues. This could be a reason why IEEE (2022) separately clarifies procedural aspects

on its homepage on how researchers should deal with found vulnerabilities in the context of conference submissions: The vulnerabilities should be reported to the manufacturer and it should be given 45–90 days to close the vulnerability before it is published. This so-called “Responsible Disclosure” procedure is a common practice in cybersecurity, not only in academia, but also in industry. We will discuss this normative practice in the following as a third approach to research ethics in the field.

Responsible Disclosure means that once a vulnerability is found in a hardware or software product, researchers should report it to its manufacturer and should grant them a certain period of time before disclosure (cf. Arora/Telang 2005, p. 20). This procedural approach is meant to strike a balance between the interests of researchers and manufacturers, but also to create pressure for patching the vulnerability in a timely fashion. At first sight, it seems to overcome the abstractness of ethical principles and to provide clear and actionable information on how to act. On closer look, however, there are many situations in which the practical issues prevail due to structural or technical problems: not all companies maintain a responsible disclosure policy and, thus, lack experience or organizational structures for handling the procedure (cf. BSI 2021, p. 71). For example, some companies simply lack corresponding points of contact, or the affected module may only be licensed from another manufacturer and not further maintained. Sometimes, there may be no manufacturer at all, but an open-source project, where volunteers may stop maintaining a critical module at any time. Besides such structural challenges, there are vulnerabilities that cannot be closed for technical reasons, e.g. because they affect hardware that cannot be fixed but is already in the hands of consumers – or because it is infeasible to access the hardware, as in the case of satellites. Here, a 90-day window does not change the ethical conflict of the situation. Hence, although Responsible Disclosure is an important tool to address dual-use issues, the procedure is not always applicable – and in such cases, we are left to the inadequate devices of abstract ethical principles.

As a fourth approach to research ethics, even though reporting vulnerabilities to manufacturers is not a silver bullet, systematic and institutionalized reports on vulnerabilities can be seen as a normatively productive tool. There are common standards, such as CVE (n. d.) (Common Vulnerability and Exposure), that serve as a collection of found vulnerabilities. Researchers (as well as other actors) can report a vulnerability to a CVE organization and after review (and possibly notifying the vendor and allowing time to develop a patch) the vulnerability will be numbered and entered into the public database. This provides a standard-

ized vulnerability identification scheme to facilitate shared communication about these vulnerabilities.

Additionally, there are also systems such as CVSS (cf. First n. d.) (Common Vulnerability Scoring System), which can be used to estimate the severity or risk level of a given vulnerability based on the evaluation of certain criteria. Such systematic collections allow getting an overview over existing vulnerabilities, which is useful from a technical perspective, but also enables national actors like the Federal Office of Information Security (*Bundesamt für Sicherheit in der Informationstechnik*, BSI) in Germany (cf. BSI 2021) or the FBI in the USA (cf. FBI 2022) to make assessments about a country's state of IT security. From a societal perspective, systematic and institutionalized reports on known vulnerabilities facilitates enforcing minimal standards, especially for operators of critical infrastructure: it may be considered due-diligence and, thus, a legal requirement to implement risk mitigations for known vulnerabilities contained here. With regard to an ethical orientation, however, the collections so far do not (yet) query enough relevant criteria from researchers to help them in their risk assessment, but they are nonetheless an interesting starting point for normative considerations.

5 Conclusions: towards a new culture of research ethics

The security of ICTs is a societal concern due to the proliferation of digital systems in our daily life. However, as we have shown, cybersecurity cannot be produced through forms of policing as cyberattacks introduce a changed risk dynamic. It has been shown that repeated non-digital attacks typically require spatial proximity, some form of direct relationship, or a linear increase in effort. Cyberattacks, in contrast, can be highly automated, globally distributed, and mutually anonymous. The resulting “economies of scale” suggest a conception of cybersecurity that focusses on the robustness of ICTs against quasi-permanent attacks that are considered an environmental fact. Therefore, cybersecurity *research* becomes an important actor in the societal production of cybersecurity – by taking on the role of attackers, but under different normative premises. As we have shown, this also gives rise to a special need for research ethics as research results can be misused by malicious actors.

The somewhat paradoxical practice of reducing societal risks by publicly demonstrating how ICTs can be attacked leads to research ethical challenges in the form of dual-use issues, especially with regard to the disclosure of security

vulnerabilities: Publication is an essential part of research, but it also increases the risk that the published results will be misused. Thus, by taking a (necessary) step in their work process, researchers may increase societal risks. This results in ethical conflicts that create the need for research ethical orientation in the field of cybersecurity.

Four approaches that respond to this need have been outlined in the last section. As has become clear, however, in many situations, the ethical challenges remain insufficiently addressed. Codes of conduct and their application at important conferences provide a general principled ethical framework that does not, however, offer sufficient orientation when dealing with vulnerabilities. As a procedural approach, Responsible Disclosure offers clear guidelines, but proves to be inadequate in more complex cases. Existing systems that allow institutionalized reporting of vulnerabilities fulfil an important societal function, but do not offer additional ethical orientation. Given the steps already taken, we believe that there is a need for further steps towards a professionalized *Ethics of Cybersecurity* that needs to be accompanied by a new research ethical culture in cybersecurity research, on the one hand, to support the research community in developing more adequate tools for orientation, but also to reflect its socio-political role in security production. It can be assumed that applied ethics and TA can make a valuable, unifying contribution here.

References

- ACM (2018): ACM Code of Ethics and Professional Conduct. <https://www.acm.org/code-of-ethics> [aufgesucht am 02.06.2023]
- Arora, A.; Telang, R. (2005): Economics of Software Vulnerability Disclosure. In: IEEE Security and Privacy Magazine 3(1), S. 20–25
- BKA – Bundeskriminalamt (2020): Cybercrime Bundeslagebild 2020. <https://www.bka.de/SharedDocs/Downloads/DE/Publikationen/JahresberichteUndLagebilder/Cybercrime/cybercrimeBundeslagebild2020.pdf> [aufgesucht am 14.08.2023]
- BMBF – Bundesministerium für Bildung und Forschung (2020): Digital, sicher und souverän in die Zukunft. https://www.bmbf.de/bmbf/de/forschung/digitale-wirtschaft-und-gesellschaft/it-sicherheit/it-sicherheit_node.html [aufgesucht am 17.05.2022]
- BSI – Bundesministerium für Sicherheit in der Informationstechnik (2021): Die Lage der IT-Sicherheit in Deutschland 2021. https://www.bmi.bund.de/SharedDocs/downloads/DE/publikationen/themen/it-digitalpolitik/bsi-lagebericht-cybersicherheit-2021.pdf?__blob=publicationFile&v=3 [aufgesucht am 14.08.2023]
- Christen, M.; Gordijn, B.; Loi, M. (Hg.) (2020): The Ethics of Cybersecurity (Bd. 21). Springer International Publishing
- CVE – Common Vulnerability Exposure (n. d.): Overview. <https://www.cve.org/About/Overview> [aufgesucht am 17.05.2022]

- Dittrich, D.; Kenneally, E., et al. (2012): The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf [aufgesucht am 17.08.2023]
- Dreißigacker, A.; von Skarczynski, B.; Wollinger, G.R. (2020): Cyberangriffe gegen Unternehmen in Deutschland. Ergebnisse einer repräsentativen Unternehmensbefragung 2018/2019. https://kfn.de/wp-content/uploads/Forschungsberichte/FB_152.pdf [aufgesucht am 17.05.2023]
- Dunn Cavely, M. (2014): Breaking the Cyber-Security Dilemma: Aligning Security Needs and Removing Vulnerabilities. In: *Science and Engineering Ethics*, 20(3), S. 701–715
- Eusgeld, I.; Fechner, B.; Salfner, F.; Walter, M.; Limbourg, P.; Zhang, L. (2008): Hardware Reliability In: Eusgeld, I.; Freiling, F.; Reussner, R. (Hg.) (2008): *Dependability Metrics: GI-Dagstuhl Research Seminar, Dagstuhl Castle, Germany*. Heidelberg/Berlin, S. 59–103
- FBI – Federal Bureau of Investigation (2022): Internet Crime Report 2022. https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf [aufgesucht am 17.05.2023]
- First (n. d.): Common Vulnerability Scoring System v3.1: Specification Document. https://www.first.org/cvss/v3-1/cvss-v31-specification_r1.pdf [aufgesucht am 17.05.2023]
- IEEE – Institute of Electrical and Electronics Engineers (2022): 43rd IEEE Symposium on Security and Privacy 2022. <https://www.ieee-security.org/TC/SP2022/cfpapers.html> [aufgesucht am 17.05.2023]
- IFIP – International Federation for Information Processing (2020): IFIP Code of Ethics and Professional Conduct. <http://ifiptc9.org/blog/2020/09/24/ifip-code-of-ethics/> [aufgesucht am 18.05.2023]
- Kováts, G.; Rónay, Z. (2023): How academic freedom is monitored: Overview of methods and procedures. [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/740228/EPRS_STU\(2023\)740228_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/740228/EPRS_STU(2023)740228_EN.pdf) [aufgesucht am 17.08.2023]
- Loschwitz, M. G. (2021): Bugs mit Vorsatz: Linux Foundation legt Analyse der Kernel-Patches vor. <https://www.heise.de/news/Bugs-mit-Vorsatz-Universitaet-legt-Bericht-zur-Analyse-der-Kernel-Patches-vor-6041701.html> [aufgesucht am 17.09.2022]
- Macnish, K.; van der Ham, J. (2020): Ethics in cybersecurity research and practice. In: *Technology in Society*, 63, 101382
- NDSS – Network and Distributed System Security (2022): Call for Papers – NDSS Symposium. <https://www.ndss-symposium.org/ndss2022/call-for-papers/> [aufgesucht am 17.05.2022]
- Reydon, T. (2013): *Wissenschaftsethik: Eine Einführung*. Ulmer
- Salter, J. (2021): Linux kernel team rejects University of Minnesota researchers’ apology. <https://arstechnica.com/gadgets/2021/04/linux-kernel-team-rejects-university-of-minnesota-researchers-apology/> [aufgesucht am 17.05.2022]
- Silomon, J. (2020): The Düsseldorf Cyber Incident. <https://ifsh.de/en/news-detail/the-duesseldorf-cyber-incident> [aufgesucht am 09.06.2023]
- USENIX – The Advanced Computing Systems Association (2021): USENIX Security Publication Model Changes. <https://www.usenix.org/conference/usenixsecurity22/publication-model-change> [aufgesucht am 17.05.2022]
- Vaughan-Nichols, S. (2021): Linux’s Technical Advisory Board reports on the UMN ‘Hypocrite Commits’ patches. <https://www.zdnet.com/article/linux-technical-advisory-board-reports-on-the-umn-hypocrite-commits-patches/> [aufgesucht am 17.05.2023]

- Wagner, M. (2020): IT-Sicherheitsforschung in rechtlicher Grauzone. Lizenz zum Hacken. In: Datenschutz und Datensicherheit – DuD 44, S. 111–120
- Weber, K.; Christen, M.; Herrmann, D. (2020): Bedrohung, Verwundbarkeit, Werte und Schaden. In: TATuP – Zeitschrift für Technikfolgenabschätzung in Theorie und Praxis 29(1), S. 11–15
- Weydner-Volkman, S.; Cassing, K. (2023): Forschende in der Angriffsrolle. Zum besonderen forschungsethischen Bedarf in der IT-Sicherheit. In: Zeitschrift für praktische Philosophie 10(1), S. 79–104

