

Operationalizing Cybersecurity Research Ethics Review: From Principles and Guidelines to Practice

Dennis Reidsma
University of Twente
d.reidsma@utwente.nl

Jeroen van der Ham
University of Twente
j.vanderham@utwente.nl

Andrea Continella
University of Twente
a.continella@utwente.nl

Abstract—Cybersecurity research involves ethics risks such as accidental privacy breaches, corruption of production services, and discovery of weaknesses in networked systems. Although literature describes these and other issues in some depth, reflection on these issues is not yet well embedded in typical Ethics Review Board procedures. In this paper, we operationalize existing guidance on cybersecurity research ethics into a proposal that can be directly implemented in an Ethics Review Board. We provide a set of self-assessment questions to effectively and efficiently probe the ethics of proposed cybersecurity research, a Coordinated Vulnerability Disclosure procedure for discoveries made in the course of research, and an outline of a university policy to institutionally embed this procedure, which could be adapted and adopted by research institutes. With this paper, we hope to contribute to more Ethics Review Boards taking up the challenge of addressing cybersecurity research ethics in an adequate and productive manner.

I. INTRODUCTION

In the field of cybersecurity, researchers as well as practitioners have to contend with ethics issues stemming from the particular nature of cybersecurity work. These ethics issues are an active topic of discussion in the field, as evidenced by the existence of conferences and workshops on this topic [20] and the presence of ethics requirements and guidelines in the call for papers of major security venues¹.

Typical topics in this field include risk of privacy breach or data leaks, risk of causing damage to production IT systems, and risk of misuse of scientific and technical results for malicious purposes (i.e., cyberattacks, mass-surveillance). A broader overview of ethical challenges in cybersecurity research is covered by recent work [11]. Furthermore, much attention has recently been paid to the question of how to deal with vulnerabilities discovered in the systems of other parties, in such a way that the discovery can be used to get the vulnerability resolved while not exacerbating the risk of others leveraging it for malicious purposes. Due to the multi-party nature of this problem, involving at least the vendor/owner of the system who may have to fix the problem, the customers/users who run the risks, and the researchers

who discover the findings, this is known as the problem of Coordinated Vulnerability Disclosure (CVD).

Regarding an industry perspective on the latter topic, in the past few years it has become common for large organizations to have explicit policies for Coordinated Vulnerability Disclosure. These policies concern on the one hand what to do when someone else discovers a vulnerability in their systems—basically, promise to responsibly investigate and fix the vulnerability and not sue the person who reported it. On the other hand the policies address what to do when they discover a vulnerability in someone else’s systems—in short, coordinate with affected parties for a responsible disclosure with a path to resolving the vulnerability for as many affected parties as possible. This may happen when a company discovers something as part of their operational work (e.g., monitoring their own network leads to detection of infection of a third party’s network), or it happens when it is part of the core business of the company to actively search out vulnerabilities (e.g., network security companies, malware scanner and firewall developers).

Also in academic work, vulnerabilities are often discovered in the course of carrying out research. Yet legal considerations are not the only aspect that should be taken into account when considering security research as many other ethical aspects should be considered. Ethics issues in academic research typically fall under the remit of an Ethics Review Board (ERB) or similar body. The central question that drives an ERB is how research should be prepared, organized, carried out, documented, and disseminated in such a way that the research does not exacerbate existing ethics risks or introduce new ethics risks, but rather prevents or mitigates them. This includes ethics risks surrounding cybersecurity research as well. However, ERBs in computer science are still in their infancy [4] and typically focus more on facets of Human Computer Interaction than on the technical computer science angle. As such, cybersecurity is under-addressed in the day-to-day practice of research ethics [17], [20].

Although certain basic principles of Cybersecurity Research Ethics have been elaborated before [10], based on theory as well as specific example cases, and although the large conferences in the field demand authors to follow certain ethical principles [20], such principles have not yet been operationalized as concrete procedures that can directly be implemented as part of an Ethics Review Board in the preparatory stage of the research. First, there are no clear guidelines on what to ask researchers in the ERB self-assessment procedure that takes place before the research is carried out, such that the right issues are probed and put on the table as input for

¹<https://www.ieee-security.org/TC/SP2023/cfpapers.html>, <https://www.usenix.org/conference/usenixsecurity23/submission-policies-and-instructions>

the conversation with the ERB reviewers. Second, there is a lack of more academia-focused best practices of ethically good cybersecurity research and procedures that researchers can directly follow. For instance, as far as we know there is no academic institute that has a policy for Coordinated Vulnerability Disclosure of discoveries made during research.

In this paper we address this gap based on our years of experience in cybersecurity research as well as in reviewing research ethics proposals in this field. After a brief introduction to the role and working of an ERB, we elaborate on the ethics of cybersecurity and cybersecurity research. We discuss some basic underlying principles and values, and we highlight the main related ethics risks and challenges that we find in literature. We then build on this to articulate three points, which summarize our contributions:

- A set of self-assessment questions that can be used to efficiently and effectively probe the research ethics of cybersecurity research as part of the ERB procedure that takes place before the research is carried out.
- A procedure for Coordinated Vulnerability Disclosure of findings made in the course of research, which will typically be carried out after completing the research.
- An outline of how such procedure can be embedded in a university policy with a proper division of responsibilities and tasks.

We release our complete questionnaire, CVD policy, and template for vulnerability notification², which can be adapted and adopted by other institutes and ERBs and can inform individual researchers about possible best practices.

Note that our goal is to provide practical and actionable guidelines with the goal of leading to better ethics practices in this research area. The questionnaire helps identify and prevent possible impact of planned research activities, and the CVD policy aims to minimize any negative impact when vulnerabilities are discovered during research activities. At the same time, we do not intend to define a set of strict rules, as ethics guidelines should be adapted to better fit local law, regulations, and research customs.

II. BRIEF OVERVIEW OF THE ROLE AND WORKING OF THE ETHICS REVIEW BOARD

An Ethics Review Board targets ethics risks in research, where such risks may either be introduced, exacerbated, mitigated, or prevented by how we prepare, carry out, document, analyze, and report the research. The European Union (EU) SATORI framework for Ethics Assessment³ puts the focus of ERBs on advice and support towards improved ethical practice, meaning that in their view the ethics review procedure is a conversation, rather than merely a gate-keeping activity—see also, for example, the European Commission (EC) guidance on ERBs in projects [16]. The heart of this review procedure is typically a self-assessment questionnaire for the researchers, which requires them to reflect on ethics risks⁴. The questions

²<https://www.utwente.nl/en/eemcs/research/ethics/coordinated-vulnerability-disclosure/>

³<https://satoriproject.eu>

⁴Which incidentally increases their awareness and understanding more generally of ethics in their research.

are ideally not meant as tick-boxes that have to be straightforwardly satisfied [15], [16]. Rather, the questionnaire probes the relevant ethics areas to get the right information on the table for starting the discussion; the answers are subsequently used as input to the review. When something remains unclear, or unresolved ethics risks are noted by the reviewers, they engage with the researchers to improve the plan until a positive opinion of the ERB can be issued.

The starting point of such a self-assessment questionnaire are big principles of research ethics such as ‘do no harm’, ‘ensure autonomy’, ‘justice’, or ‘equal benefit’. However, these big principles are hard for researchers to work with directly; asking the question “*Is autonomy guaranteed in your research?*” is not likely to get information on all pertinent ethics risks of the project on the table. Therefore, the big principles are typically operationalized into a limited set of specific self-assessment questions that are easier to work with and more conducive to eliciting the salient information for a specific type of research.

Questions in the self-assessment procedure work well if they efficiently (i.e., with a limited set of questions [16]) and effectively probe the right information. This allows the reviewers to quickly get to the heart of issues, to ask for more pointed elaborations on specific points, and to give more concrete and actionable advice—without the researcher first having to read up on decades of research ethics literature to understand what they should explain in their self-assessment. Rather than generically asking “*How do you ensure voluntariness in your plan?*” the questionnaire typically asks more concrete things such as “*What are your consent forms and information brochures; do your human subjects include persons from vulnerable populations that may have limited capacity to consent; and are there power relations that may threaten voluntariness of participation?*”. Although these operationalized questions do not always cover every possible issue, they have proven in practice to be effective in probing enough of the risks and ethics issues such that the reviewers feel confident in providing advice on the basis of the answers—especially in combination with incidental help from an ethics philosopher whom the SATORI framework advises to have on board in the ERB. Ideally, the resulting advice and requests-for-revision from the ERB to the researcher are actionable and concrete, and thus helpful in improving research practice. For example, instead of saying “*Please explain better how you guarantee justice*” a reviewer might advise “*Your sampling is unbalanced; this means your technology will be tailored to a too limited demographic; please show how you can improve or mitigate this*”. Such advice is based on a long-standing experience with best practices of ethical research in a field.

Of special note is the ‘first’ question in the questionnaire, which roughly comes to “do questions of research ethics come up for your research?”. While on the surface this seems a trivial question, it is not always so. As an example, in 2021 a research paper was initially accepted and then retained in the program of the IEEE Symposium on Security and Privacy (S&P)⁵. The presented research [19] involved the introduction of vulnerabilities into open-source software projects through minor patches, with the goal of studying whether such minor

⁵https://www.ieee-security.org/TC/SP2021/downloads/2021_PC_Statement.pdf

patches could lead to new, hard-to-detect bugs. While initially, the researcher’s ERB could not adequately support an ethics review of their work, the deceptive nature of the study, the lack of debriefing of the study participants involved, and the risk of inadvertently introducing bugs into critical open-source code—with an imbalance of risks over benefits—made the authors and the PC Chairs realize that this research did not meet the ethics criteria discussed in the Menlo Report [1]

As visible in the examples above, for research in Human Computer Interaction there is quite some experience in posing the right probing questions and translating the answers to actionable advice based on best practices. For cybersecurity research this is not yet the case. Therefore in the remainder of this paper we work towards operationalizing the available general insights on cybersecurity research ethics into concrete best practices and procedures that can be directly implemented in the ERB review process.

III. ETHICS PRINCIPLES AND VALUES IN CYBERSECURITY

Although ERBs do not yet have a good working practice for the research ethics of cybersecurity, earlier work has outlined several underlying principles that could be used as a starting point [2], [6], [10], [11], [17]. These papers are based on a mixture of discussing and analyzing past cases of problematic research, pragmatic overviews on what was already out there in terms of cybersecurity research ethics support, and knowledge from ethical theories.

The first main work in this domain is the Menlo Report [1], which adapted the earlier Belmont Report [12] for information and communication technology research. Despite being an important reference, the Menlo Report does not cover and discuss all potential risks with cybersecurity research, such as the risk of incidental findings. Instead, a more recent work [10] provides a more comprehensive set of important ethics issues in this domain. We gathered risks from literature, especially [3], [10], and summarize such issues, together with their explanations and related principles, in Table I. Main organizing structure is the little ‘scenarios’ that researchers might encounter in their work and can easily relate to, while for each we provide short explanation, and link to the main underlying ethics principles.

While the literature provides us with an overview of the “risk scenarios”, the challenge, though, is to turn these risk scenarios into operational practice in the day-to-day functioning of the ERB procedure. Unfortunately, it is not always clear how exactly these risks apply to specific research works, as their explanation is often too broad or abstract. As mentioned before in previous sections, operationalizing ethics review involves framing best practices that the researcher can easily relate to, and self-assessment questionnaires they can meaningfully answer without reading the whole literature up-front.

We now first discuss our self assessment questionnaire that we developed to elicit the right information in ERB procedures. Then, we discuss the best practice that we developed specifically for issues of Coordinated Vulnerability Disclosure (CVD) that may come up in research.

IV. REVIEWING THE ETHICS OF CYBERSECURITY RESEARCH IN AN ERB

The first contribution offered in this paper is an (initial) set of probing questions to be used in the self-assessment. These questions are part of a larger questionnaire addressing issues such as human participants, use of databases, and other common topics; here we only focus on the questions regarding cybersecurity ethics. Which questions will most easily elicit the right issues, without requiring the researcher to know all the examples or have deep ethics knowledge, and without requiring them to answer a too complex set of questions, given that we know what are the issues that should be on the table? The set of questions should be clear and concise, firstly because a typical researcher will not read related papers [10] merely to know how they should approach their ethics, and secondly because long and onerous questionnaires are a reason for people to disengage from the ethics review process and to approach it as a tick-box exercise in which the goal is not to ethically reflect but to “provide the answers that get one past the ERB as quick as possible”. Despite conciseness, the questions should allow reviewers to effectively get at the real issues.

Based on existing literature and our own experience of the most fruitful discussions that we have with our researchers, we constructed a set of questions that provide the reviewer of the ERB with salient information in two ways. The direct answer is sometimes indicative of possible issues, plus sometimes certain answers are a signpost for followup questions by the reviewer. The number of questions is kept as small as possible, to not scare away the researcher – so we do not elaborately ask about every single risk separately. Rather, the single questions are each indicative of multiple risks at the same time. On the other hand, also most ethics risks in Table I are directly or indirectly addressed by more than one question, making it more likely that the relevance of that risk is brought to the reviewer’s attention one way or another. For example, Q4 directly asks for accidental discovery, but if researcher says “no” because they do not know they should have said “yes”, then the answer to Q3 may still hint at a possible risk of discoveries (because the research is being done “out in the real Internet”). This, then, helps reviewer re-assess the answer to Q4 and look for a followup to that issue in Q5. Similarly, Q5 obviously is about R5, but also Q1 and Q3 may have a followup about R5 that may either trigger extra answers under Q5, or may be a cause for the reviewer to look closely when this risk does not come up under Q5. And if production systems are involved, it becomes more relevant for the reviewer to keep an eye out for possibly unaddressed issues of accidental privacy breach, and to take a look at possible issues of consent. As such, the questions are used as a primer for researchers and reviewers to elicit possible impact and dilemma’s; the expertise of the ERB reviewer in seeing the implications of the answers is a large factor in the possible effectiveness of these questions. In the rest of this section we present the questions that we have, and how they relate to various risks.

OPENING QUESTION: *Will the research involve any cybersecurity or online privacy issues, such as the possible discovery of security vulnerabilities, experiments with malicious software (e.g., computer viruses), or the discovery and investigation of illegal activities on the Internet?* – Used to help people identify whether they should fill out this particular set of questions.

TABLE I. TYPICAL ETHICS RISKS IN CYBERSECURITY RESEARCH, MOSTLY DERIVED FROM [10] EXCEPT R4 AND R9 WHICH WERE TAKEN FROM [3]. EACH RISK IS EXPLAINED THROUGH A LITTLE SCENARIO-LIKE DESCRIPTION THAT A RESEARCHER CAN EASILY RELATE TO THEIR DAY-TO-DAY RESEARCH PRACTICE, AS WELL AS CONNECTED TO THE LARGER UNDERLYING ETHICS PRINCIPLES. NOTE THAT THE FIELD OF ETHICS IS CHANGING QUICKLY AND SO NEW RISKS MIGHT EMERGE AND REQUIRE ATTENTION IN THE FUTURE.

#	Risk	Explanation	Principle
R1	Accidental discovery	E.g., discovering evidence of criminal activities while you were not searching for it. Do you report this, or keep it confidential? Not just a legal (GDPR or other) matter but also concerns a moral right to privacy.	Privacy; Consent
R2	Publication of private data	Accidentally disclosing private data of people, to the researchers or to the outside world, that was accessed deliberately as part of the research, or as a side effect of accessing other information. Concerns legal (GDPR or other) and moral right to privacy.	Privacy; Consent
R3	Damaging production systems	E.g., flooding and flooring networks with live measurements; accidentally destroying a server; accidentally releasing harmful software into the wild.	No Harm
R4	Misuse by malicious actors	E.g., knowledge on hacking, vulnerabilities, viruses, and similar can also be used by malicious actors to intentionally carry out harmful activities.	Misuse; Dual Use (Military Applications)
R5	Discovering vulnerabilities	Knowing about a vulnerability means the question arises to disclose or not, and how to approach that.	No Harm
R6	Reprisal against researcher	Not all vendors respond reasonably when vulnerabilities in their systems are reported to them – there is a wide range of risks, including putting the researcher under financial, legal, reputational, or other pressure.	Researchers’ well-being
R7	Lack of consent	In cybersecurity research it is not always obvious how to obtain consent and whom to ask for it. Whose machines are addressed in the research? In whose environment are you operating?	Consent
R8	Damaging (reputation of) people or companies	E.g., generating “suspicious” behavior from someone else’s machine; publishing weaknesses in a company’s systems which can have commercial impact on that company.	No Harm; Consent
R9	Illegal behavior by researcher	Some research requires the researcher to carry out activities that are <i>de facto</i> illegal, like installing malicious code in systems.	Consent; Deception; Illegality; Researchers’ well-being

People who know they should, can easily answer this one, for others it provides a compact hint of reasons why to fill out this section. We find that the default assumption of inexperienced researchers will be “no, not relevant for my research” but that adding a small handful of hints here might increase the number of researchers filling out these questions.

Q1 SECURITY WEAKNESSES: *Could your research result in the identification of security weaknesses in existing systems?* – R4, R5, R6, but sometimes also R1, R2

Q2 MALICIOUS SOFTWARE: *Will your research involve experiments on malicious software (e.g., computer viruses) or real-world attacks (e.g., denial of service attacks)?* – R3, R4, R7, R9

Q3 INTERNET: *Will your research involve external machines on the Internet?* – R3, R1, R2, but also R5, R7, R8, R9

Q4 ACCIDENTAL DISCOVERY: *Might your research lead to the accidental discovery of illegal behavior or behavior that could pose a risk to others, either directly or indirectly, on the Internet?* – R1, R6, R7, R8

Q5 COORDINATED VULNERABILITY DISCLOSURE: *Please explain concisely, and in relation to your answers above, how you will deal with the various potential issues raised in the previous questions of this section and how you*

will follow the UT procedure for Coordinated Vulnerability Disclosure. – R5, but also R1, R2, R6; furthermore, “the various potential issues” is a catch-all inviting the researcher to reflect on all risks.

V. CYBERSECURITY AND COORDINATED VULNERABILITY DISCLOSURE IN RESEARCH

The second contribution offered in this paper concerns a specific best practice way of working for one facet of the ethics of cybersecurity, namely Coordinated Vulnerability Disclosure of discoveries made as part of research. As far as we know, there is no academic institute yet that has such a policy at institutional level. We developed a protocol and policy that we are currently implementing at our university to help our researchers carry out these (sometimes complex) procedures.

A. A Brief History of Vulnerability Disclosure

Ever since the discovery of the first security vulnerability in software, there have been discussions on how to disclose these and to whom. This discussion boiled down to three different approaches for vulnerability disclosure: full disclosure, coordinated disclosure or no disclosure⁶. These discussions were later influenced by rewards and bug bounties.

⁶<https://addxorrol.blogspot.com/2019/08/rashomon-of-disclosure.html>

With *full disclosure* the researcher reports a vulnerability to the public immediately. This style of disclosure was prevalent in the early years, when the common thinking was that software vendors were not trying hard enough. So they should not be helped to improve their security, and instead everyone should learn about these vulnerabilities at the same time.

In *coordinated disclosure* a researcher coordinates with the owner or vendor of the software to disclose the vulnerability. Initially the researcher only discloses the vulnerability to the owner or vendor so that they are in a position to create a remediation or patch. When this is available, the researcher and vendor can coordinate to disclose the vulnerability to the wider public. Coordinated vulnerability disclosure is the prevalent name for this style currently. In earlier years the term *responsible disclosure* was used. This term came into disfavor as it can be seen as to put the responsibility of this action with the researcher who is trying to improve security by reporting the vulnerability to the right parties. However, the owner or vendor of the affected system bears responsibility of this as well, as they are the one who introduced the vulnerability, and are most likely in the best position to fix the vulnerability. So, the adjective ‘coordinated’ is preferred over ‘responsible’ to indicate that both parties are equally responsible in the process.

The third option is *no disclosure*, the researcher does not disclose the vulnerability at all, or at least not to the owner or vendor. They can use the vulnerability for their own activities, or they can sell the vulnerability to others, for private use.

Bug bounties and rewards came into the discussion later, when security researchers felt being taken advantage of. With the movement ‘No More Free Bugs’ many security researchers stopped reporting security vulnerabilities to vendors, demanding better compensation for their time. This turned out to give the vendors some breathing room to fix outstanding security vulnerabilities, while discussions on better compensation and recognition continued. Eventually many vendors started rewarding researchers for disclosing vulnerabilities, either in their regular vulnerability disclosure program, or through bug bounty platforms, such as HackerOne or Bugcrowd.

Last, but certainly not least in the vulnerability disclosure discussion is the legal aspects of vulnerability disclosure. Another part of the ‘No More Free Bugs’ discussion was that a large part of the security research activities for finding security vulnerabilities may be breaking the law, and putting the researchers at risk⁷. This led to the development of vulnerability disclosure policies to help researchers disclose to organizations with vulnerabilities. Over the years there has been numerous developments in this context, supporting organizational processes [7]–[9], and support from international organizations [5], [13], [14]. Such a development has also been supported and carried by some countries and governments. For instance, in The Netherlands there has been an official legal policy on vulnerability disclosure since 2013. The Public Prosecutor has published a position paper⁸ on how they deal with “ethical hacking”, meaning that the court tests for general interest, proportionality and subsidiarity.

⁷<https://web.archive.org/web/20111129105008/http://trailofbits.com/2009/03/22/no-more-free-bugs/>

⁸<https://www.om.nl/documenten/richtlijnen/2020/december/14/om-beleidsbrief-ethisch-hacken>

B. Researcher Disclosure Policy

The traditional Coordinated Vulnerability Disclosure policy is published by organizations outlining expected behavior of researchers, and aiming to provide ‘safe harbor’ assurances. Similarly, bug bounty platforms have policies describing expected researcher behavior.

There have been only a few initiatives publishing policies from a researcher perspective describing the process and timelines that are written from a researcher/discloser point of view. The first policy has been published by Google Project Zero⁹. Most notably this led to a discussion on the 90 day deadline for disclosure.

Google Project Zero set a strict 90 day deadline for disclosure processes. This meant that 90 days after their first contact with a receiving organization, the details of the vulnerability would be published, regardless of the response of the receiving organization. Until then, a period of 90 days was taken as a suggestion, but was not strictly applied by researchers. Many software vendors felt threatened by this strict deadline, but eventually this turned out to be a workable period. This also led to CISA (the Cybersecurity & Infrastructure Security Agency) setting a 45 day deadline in case ‘a vendor is unresponsive, or will not establish a reasonable timeline for remediation’¹⁰.

The researcher disclosure policy made it clear beforehand how the discoverers would report vulnerabilities, how they would go about finding contact information, and what they would expect from receiving organizations. These guidelines prevent discussions while handling incidents under stress.

C. Our Procedure

Our procedure provides researchers and students with clear guidelines for conducting vulnerability discovery activities and coordinating with vendors to disclose and mitigate the discovered vulnerabilities. Our intent is to disclose vulnerabilities in the most helpful way to the community by ensuring confidentiality during the process, working with affected parties to find and test fixes, and aiming to inform all the impacted entities so that they can protect themselves by deploying patches and updating their systems. Because the disclosure process can be complex and long, we require supervisors to be responsible for conducting the procedure for the vulnerabilities reported to them by their students. Besides, researchers are asked to keep a record of the communications concerning the CVD.

First, researchers should identify the right contact for reporting a vulnerability. Contact methods could include but are not limited to using the contact information in the CVD policy of the owner or vendor, the ‘security.txt’ contact information [7], emailing security reporting emails (security@ or secure@), filing bugs without confidential details in bug trackers, or filing support tickets. Once contacts have been identified, the researchers should send a first notification, including that the vulnerability was found in a scientific environment, proposing a deadline for publication of the reported issue to prevent deadlock because of no response, and stating that they are willing to negotiate publication date, pending response and remediation actions. Importantly, we expect researchers to

⁹<https://googleprojectzero.blogspot.com/p/vulnerability-disclosure-faq.html>

¹⁰<https://www.cisa.gov/coordinated-vulnerability-disclosure-process>

write such a notification in a friendly and open tone. For this, we provide researchers with template for notifications¹¹

In case of no response, our procedure includes sending reminders after 21 and 60 days. Additionally, in these cases, researchers should attempt to contact software distributors. For instance, in case of vulnerabilities found in an Android app present in the Google Play Store, the researchers should contact Google.

If no fix is available at the end of the agreed publication date (e.g., after 90 days), the researchers should notify the contact of the intent to disclose the reported issue—in case of mitigating circumstances, it is possible to extend the deadline. When either the issue is fixed or the (extended) deadline is expired, the researchers can disclose the vulnerability. Depending on the nature of the problem, there may be a few disclosure paths: 1) disclose the vulnerability publicly, 2) disclose it directly to the people using the project, or 3) issue a limited disclosure first, followed by a full public disclosure. The path to follow is determined on a case-by-case basis.

Along the whole disclosure process, our procedure involves documentation of the performed steps and timeline. We also publicly publish our procedure on the university website to make affected vendors and parties aware of our policy.

In the event that a vulnerability disclosure turns out to be very complicated or difficult, we reach out to the governmental CERT, NCSC-NL¹². They can help in finding the right contacts, provide support for multi-party vulnerability disclosure cases, or help as a sparring partner from their extensive experience in performing vulnerability disclosure.

D. Embedding the CVD Procedure in University Policy

The procedure described in the previous section has been in *de facto* operation for close to a year now. Currently, we are working with the Strategy & Policy department and the Executive Board of the University to make the procedure part of an official University Policy.

The University Policy will make it clear to researchers how they are expected to behave, and will give the University somewhat more formal leverage in demanding that researchers actually follow these procedures. At the same time, it provides researchers with assurance that they will be protected if they follow the guidelines. With the University publishing this as an explicit policy, the Executive Board are also committing themselves into following this policy and protecting their researchers from legal consequences.

At the same time, these procedures make it clear to recipients of disclosure notices how the University will handle this process, and what can be expected and when. Publishing this policy generally beforehand can prevent many discussions during the disclosure process, and will hopefully also prevent the recipient from feeling like they are being pressured or even extorted [18].

¹¹<https://www.utwente.nl/en/eemcs/research/ethics/coordinated-vulnerability-disclosure/>

¹²<https://www.ncsc.nl/>

VI. DISCUSSION AND FUTURE WORK

The self-assessment questionnaire and the procedure for Coordinated Vulnerability Disclosure have been actively used in our own Ethics Review Board for some time now. Although the number of research protocols submitted for ethics review on this topic is a fraction of what our ERB receives for human subject based research, the cases that we do get often lead to interesting discussions with the researchers, potentially leading to improvement of research practice – although the latter still needs to be investigated. Furthermore, we have seen the first actual cases of research-based CVD procedures come up in our practice. One of the challenges that we see revolves around the scalability of the procedure since some research outcomes appeared to affect tens of thousands, rather than a handful, of systems, vendors, and end users. The future will show how best practices in disclosure procedures will have to evolve to take this and other challenges better into account.

Clearly, future work includes iterating over the questionnaire to make it easier to use, keeping it up to date with technical developments that may change the ethics risks or introduce new issues, and teaching such topics to students. Furthermore, the type of research addressed here is the only source of ethics issues. Cybersecurity research also relates in specific ways to human subject based research, artificial intelligence research, database research, and other topics – often leading to a particular take on the known ethics issues in those other domains. Thus, future work should explore whether other “modules” of best practices are needed for other classes of ethical problems that security researchers may need to address during or after the start of their research.

Our work also points the way for institutes to organize their own ethics support for cybersecurity research by adapting our guidelines. For any ERB who wants to implement similar procedures and self assessment questionnaire, we want to reiterate the importance of having an ethics philosophy specialist on board as advised in the Satori Framework¹³ and for any University Executive Board who wants to adopt our suggested policy, we want to highlight the importance of offering full support to researchers who follow these procedures, in the (relatively rare) cases where the procedures are not straightforward smooth sailing. Explicitly expressing this support is, we feel, a major factor in the researchers feeling intrinsically motivated to engage with this ethics process.

VII. CONCLUSION

Regarding the ethics of Cybersecurity and Coordinated Vulnerability Disclosure in academic research, we provided a best practice, a sketch towards a university policy, and a partial self-assessment questionnaire that may help researchers reflect on risks and mitigations and help the Ethics Review Board elicit salient information to carry out their reviewing and advisory tasks. These can be adopted by ERBs and institutes, possibly adapted to better fit local law, regulations, and research customs. We hope this leads to better practices in this type of research, and that it will help academia to better take into account, and deal with, ethics issues in this field.

¹³<https://satoriproject.eu/framework/section-4-ethics-assessment-procedures/>

ACKNOWLEDGMENT

We would like to thank our reviewers for their valuable comments. We would also like to thank the people who supplied us with relevant research cases to review; the people who supplied actual cases of Coordinated Vulnerability Disclosure to handle; the people who helped improve the self assessment ethics checklist; and the people who are helping us improve the policy land in our University.

REFERENCES

- [1] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The Menlo report," *IEEE Security & Privacy*, vol. 10, no. 2, pp. 71–75, 2012.
- [2] R. Bhalerao, V. Hamilton, A. McDonald, E. M. Redmiles, and A. Strohmer, "Ethical practices for security research with at-risk populations," in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2022, pp. 546–553.
- [3] M. Botes and G. Lenzini, "When cryptographic ransomware poses cyber threats: Ethical challenges and proposed safeguards for cybersecurity researchers," in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2022, pp. 562–568.
- [4] P. Brey, "Research ethics guidelines for the engineering sciences and computer and information sciences," in *Codes of Ethics and Ethical Guidelines*. Springer, 2022, pp. 15–34.
- [5] European Union, "Proposal for a directive of the european parliament and of the council on measures for a high common level of cybersecurity across the union, repealing directive (eu) 2016/1148 (nis2 directive)," 2022. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2020%3A823%3AFIN>
- [6] A. Eusebi, M. Vasek, E. Cockbain, and E. Mariconti, "The ethics of going deep: Challenges in machine learning for sensitive security domains," in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2022, pp. 533–537.
- [7] E. Foudil and Y. Shafranovich, "A File Format to Aid in Security Vulnerability Disclosure," RFC 9116, Apr. 2022. [Online]. Available: <https://www.rfc-editor.org/info/rfc9116>
- [8] ISO, "Information technology - security techniques - vulnerability disclosure," ISO, Tech. Rep. ISO/IEC 29147, 2018.
- [9] —, "Information technology - security techniques - vulnerability disclosure," ISO, Tech. Rep. ISO/IEC 30111, 2018.
- [10] K. Macnish and J. van der Ham, "Ethics in cybersecurity research and practice," *Technology in Society*, vol. 63, p. 101382, Nov 2020.
- [11] —, "Ethical approaches to cybersecurity," in *The Oxford Handbook of Digital Ethics*, C. Véliz, Ed. Oxford University Press, Mar 2022.
- [12] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research of the United States, *The Belmont report: ethical principles and guidelines for the protection of human subjects of research*. Department of Health, Education, and Welfare, National Commission for the . . . , 1978, vol. 2.
- [13] OECD, "Understanding the digital security of products," *OECD Digital Economy Papers*, Feb 2021. [Online]. Available: <http://dx.doi.org/10.1787/abea0b69-en>
- [14] 2022. [Online]. Available: https://elearning.osce.org/courses/course-v1:OSCE+TNTD-CYBERCVD+2022_04/about
- [15] B. Stahl, J. Timmermans, S. Rainey, and M. Shaw, *Ethics in innovation management as meta-responsibility*. Routledge, 02 2019, pp. 435–456.
- [16] N. Stingelin, G. V. Steendam, J. Rath, J. A. Houghton, J. Schmucker von Koch, and R. Iphofen, "The roles and functions of ethics advisors/ethics advisory boards in EC funded projects," 2021. [Online]. Available: https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/guidance/roles-and-functions-of-ethics-advisory-ethics-advisory-boards-in-ec-funded-projects_he_en.pdf
- [17] J. van der Ham and R. van Rijswijk-Deij, "Ethics and internet measurements," *Journal of Cyber Security and Mobility, Volume 5, Issue 4*, oct 2017. [Online]. Available: <http://dx.doi.org/10.13052/jcsm2245-1439.543>
- [18] T. Walshe and A. Simpson, "Coordinated vulnerability disclosure programme effectiveness: Issues and recommendations," *Computers & Security*, vol. 123, p. 102936, 2022.
- [19] Q. Wu and K. Lu, "On the feasibility of stealthily introducing vulnerabilities in open-source software via hypocrite commits," in *Proc. Oakland*, 2021.
- [20] Y. Zhang, M. Liu, M. Zhang, C. Lu, and H. Duan, "Ethics in security research: Visions, reality, and paths forward," in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2022, pp. 538–545.