



Practical Cybersecurity Ethics: Mapping CyBOK to Ethical Concerns

Ivan Flechais
University of Oxford
United Kingdom
ivan.flechais@cs.ox.ac.uk

George Chalhoub
University College London
& University of Oxford
United Kingdom
g.chalhoub@ucl.ac.uk

ABSTRACT

Research into the ethics of cybersecurity is an established and growing topic of investigation, however the translation of this research into practice is lacking: there exists a small number of professional codes of ethics or codes of practice in cybersecurity, e.g. the ISSA or the UK Cyber Security Council's code of ethics, however these are very broad and do not offer much insight into the ethical dilemmas that can be faced while performing specific cybersecurity activities. In order to address this gap, we leverage ongoing work on the Cyber Security Body of Knowledge (CyBOK) to help elicit and document the responsibilities and ethics of the profession.

Based on a review of the existing literature on the ethics of cybersecurity, we use CyBOK to frame the exploration of ethical challenges in the cybersecurity profession through a series of 15 interviews with cybersecurity experts. Our approach is qualitative and exploratory, aiming to answer the research question "What ethical challenges, insights, and solutions arise in different areas of cybersecurity?". Our findings indicate that there are broad ethical challenges across the whole of cybersecurity, but also that different areas of cybersecurity can face specific ethical considerations for which more detailed guidance can help professionals in those areas. In particular, our findings indicate that security decision-making is expected of all security professionals, but that this requires them to balance a complex mix of different technical, objective and subjective points of view, and that resolving conflicts raises challenging ethical dilemmas. We highlight our participants' concerns about the growing use of AI technology in cybersecurity, and discuss the implications of applying AI to decision-making.

We conclude that more work is needed to explore, map, and integrate ethical considerations into cybersecurity practice; the urgent need to conduct further research into the ethics of cybersecurity AI; and highlight the importance of this work for individuals and professional bodies who seek to develop and mature the cybersecurity profession in a responsible manner.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Security and privacy** → *Usability in security and privacy*; Social aspects of security and privacy.

KEYWORDS

cyber security, ethics, cybok, security

ACM Reference Format:

Ivan Flechais and George Chalhoub. 2023. Practical Cybersecurity Ethics: Mapping CyBOK to Ethical Concerns. In *New Security Paradigms Workshop (NSPW '23)*, September 18–21, 2023, Segovia, Spain. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3633500.3633505>

1 INTRODUCTION

Cybersecurity professionals wield enormous power. In pursuing the goal of protecting computer systems, their responsibilities can variously lead them to monitor people and review sensitive information; investigate threat actors; document and prosecute insiders; subject others to attack through penetration tests or "ethical hacking" activities; decide on limiting, quarantining, revoking and denying access to systems and data in the face of ongoing attacks; or deal with the anguish, betrayal, and trauma arising from harmful cyber attacks. As noted by Christen et al. [14], "*Overemphasising cybersecurity may violate fundamental values such as equality, fairness, freedom or privacy. However, neglecting cybersecurity could undermine citizens' trust and confidence in the digital infrastructure, in policy makers and in state authorities.*" Such power over other people's actions and freedoms should come with clear, transparent, and detailed ethical oversight, however this is far from the case in practice.

Research into the ethics of cybersecurity is an established and growing topic of investigation (see Sections 2.1 and 2.2), however the translation of this research into professional practice is lacking: there exists a small number of professional codes of ethics or codes of practice in cybersecurity, e.g. the ISSA or the UK Cyber Security Council's code of ethics, however these are very broad and do not offer much insight into the ethical dilemmas that can be faced while performing specific cybersecurity activities.

Significant efforts are underway to improve the maturity of cybersecurity: ranging from improvements in secure software development lifecycles, data protection law and regulation, cyber insurance, or efforts to codify the Cyber Security Body of Knowledge (CyBOK). In order to investigate the gap we identify between the research and practice of ethics in the cybersecurity profession, we propose the following research question "*What ethical challenges, insights, and solutions arise in different areas of cybersecurity?*" Our approach is qualitative and exploratory: drawing on CyBOK and



This work is licensed under a Creative Commons Attribution International 4.0 License.

NSPW '23, September 18–21, 2023, Segovia, Spain
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1620-1/23/09.
<https://doi.org/10.1145/3633500.3633505>

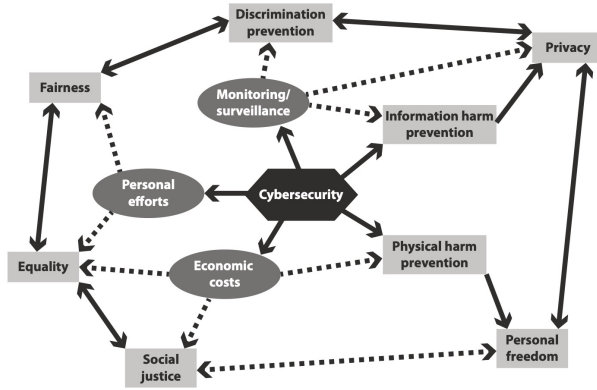


Figure 1: Value Conflicts in CyberSecurity. Arrows with continuous lines show positive (i.e., supporting) relations, whereas arrows with dotted lines show conflicting relations. Reproduced from [15] through CC-BY license

the principlist framework for cybersecurity ethics by Formosa et al. [25], we engage with professionals in the field to elicit, map, and deepen our understanding of the key issues in specific areas of cybersecurity.

In Section 2, we review the current state of cybersecurity ethics in both research and professional practice. Section 3 describes our research approach, and in Section 4 we present our results. We discuss these in Section 5, before concluding with suggestions for future work.

2 BACKGROUND

The ethics of cybersecurity is establishing itself as a field of ethical research in its own right, sharing similarities with other fields such as bioethics, digital ethics, or the ethics of artificial intelligence, but having its own specific characteristics. While the ethics of cybersecurity has been dominated by concerns around privacy values, it is becoming clear that there are wider ethical challenges that arise [9–13, 20, 53]. Christen et al. [15] explore these and Figure 1 summarises some of the key relationships between cybersecurity values and other ethical values.

In the following we start by presenting an overview of the underlying principles of cybersecurity ethics. We then explore different ethical frameworks for cybersecurity which aim to provide a more pragmatic structure to help evaluate different ethical questions. Then we provide a brief outline of the Cyber Security Body of Knowledge (CyBOK), and conclude by focusing on how professional codes of practice currently provide guidance for practitioners.

2.1 CyberSecurity Ethical Principles

There is a wide variety of different approaches looking at the ethics of cybersecurity [35]. Aiming to categorise these, Macnish and van der Ham [38] argue that there are three broad approaches, calling the first “bottom-up” which examines detailed cases and identifies ethical issues arising from these. In contrast, the second is labelled “top-down” and focuses on ethical values or theories as a starting point which are then applied to the cybersecurity

context. Macnish et. al. also identify a third approach which they call “pragmatist”, which focuses on the practices of cyber security professionals. While they note that such approaches usually focus more on framing the values of information security (such as Confidentiality, Integrity, Availability) rather than identifying ethical issues, principles, or solutions, it is important to note that there are a variety of professional codes of practice that fall into this category (see Section 2.4). Formosa et al. [25] argue for two broad categories instead. The first aims to apply established ethical theories (consequentialism, deontological ethics, and virtue ethics [42]) for which the ethical textbook by Manjikian [39] is a prominent example. The second category aims to outline a series of mid-level and domain-specific principles, an approach known as “principlism”. In both categories, Formosa et al. [25] note that casuistry (an applied ethics approach that uses case-based reasoning to derive ethical insights) is widely used.

Examining case studies of moral or ethical dilemmas in cybersecurity is a widespread approach that provides context and insight. Common case studies include generic example applications such as penetration testing or ethical hacking [34], encryption in the context of privacy vs state surveillance (e.g. Rogaway [48] whose opening statement “*Cryptography rearranges power*” neatly encapsulates a core issue), or electronic voting [47]. In contrast, other case studies focus on highly detailed and specific case studies, an example of which is the examination of the ENCORE project by Byers [7]. The ENCORE project [6] proposed a method to explore online censorship by harnessing cross-origin requests to covertly induce web browsers running on computers in various different countries into contacting specific websites and reporting back. Given (1) that this activity was non-consensual for the owners of those computers, (2) that there is significant potential harm to these owners arising from repressive regimes tracking attempts at accessing censored material, and (3) that the authors had already deployed and tested their approach in the wild, this research and its publication led to a significant ethical debate in the academic community. A key aspect of this debate centred around the fact that the research was approved by the authors’ Institutional Review Board, who did not identify ethical issues arising from the technical details of the research. Following lengthy deliberation, the Program Committee for SIGCOMM voted to publish the paper, however they took the unprecedented decision to document their ethical concerns in a statement at the top of the paper [7].

Whether drawing from case studies or applying ethical theories to the cybersecurity context, a key aim has been to identify and specify the core ethical values of cybersecurity. The chapter by van de Poel [57] provides a helpful overview of some of these, noting that ethical cybersecurity values can be clustered into aspects of security, privacy, fairness, and accountability. Christen et al. note in their introductory chapter that cybersecurity involves a balance between fundamental values such as equality, fairness, freedom or privacy and the need for protecting citizens’ trust and confidence in the digital infrastructure, in policy makers, and in state authorities [14].

2.2 CyberSecurity Ethical Frameworks

Ethical frameworks aim to provide structure to help evaluate ethical questions. Loi and Christen [37] provide a helpful overview of these and note that there are several notable ethical frameworks in cybersecurity.

2.2.1 Human Rights Frameworks. The first type of framework is based on human rights and how these are embedded in various legal and regulatory frameworks. Hildebrandt [32] discusses how rights can interact with cybersecurity, focussing specifically on privacy, data protection, non-discrimination, due process and free speech. Under EU law, these rights are protected, however cybersecurity activities (such as monitoring, profiling, filtering content) can come in conflict with these rights. Hildebrandt [32] argues that in the case where a conflict is resolved through a trade-off, infringing measures have to be balanced by effective safeguards. To achieve this balance, Hildebrandt draws on the triple test, derived from the second paragraph of Art. 8 of the European Convention of Human Rights, which requires that a right's infringement "*must be in accordance with the law, necessary in a democratic society and have a legitimate aim.*" While a detailed review of various legal frameworks is beyond the scope of this paper, it is notable that ethical concerns have played a significant role in the formulation of EU data protection regulation, such as the General Data Protection Regulation, which enshrines seven principles (Lawfulness, fairness and transparency; Purpose limitation; Data minimisation; Accuracy; Storage limitation; Integrity and confidentiality (security); Accountability) and eight individual rights (The right to be informed; The right of access; The right to rectification; The right to erasure; The right to restrict processing; The right to data portability; The right to object; Rights in relation to automated decision-making and profiling).

2.2.2 Theory of Contextual Integrity. The second approach is Nissenbaum's theory of Contextual Integrity [43]. This theory has enjoyed widespread success arising from its characterisation of privacy violations as violations of social norms from the transmission of information between persons. Social norms are grounded in the specific contexts of each situation, allowing the exploration of privacy in a manner that is sensitive to societal, cultural, and wider contextual factors.

2.2.3 Ethics of Risk. The third type of framework explores the ethics of risk. Drawing on the significant body of work from Sven Ove Hansson on the ethics of risk [29], Macnish and van der Ham [38] argue that cybersecurity is the inverse of risk. Where risk is defined as the likelihood of harm arising from a threat, security grows as risk reduces; conversely, as security decreases, the likelihood and impact of harm increases. In positioning cybersecurity in terms of risk, a number of interesting observations can be made. The first observation is that security is always forward-looking: since risk aims to anticipate the likelihood and impact of future threats, then likewise security is only concerned about the future – should an incident occur, it is no longer in the realm of risk or security but in the realm of harm, crisis and security failures. The second observation was made by Herington [31] who noted that security has subjective, objective, and affective dimensions: a person could be secure (objective), believe that they are secure (subjective), but not feel secure (affective). These qualities are interrelated in that they

can influence one another, however they are distinct and security needs to satisfy a number of potentially conflicting perspectives: rational (objective), personal (subjective), and emotional (affective).

Sven Ove Hansson and [29], Macnish and van der Ham [38] also highlight that there are four notable ethical implications from looking at risk: (1) the distinction between objective and subjective harms, (2) the challenges of calculating probabilities, (3) the recognition of fallacies, and (4) the problems arising from risk thresholds and distribution. Evaluating harm (1) highlights the inherent quality of security having subjective, objective and affective facets, however most approaches tend to aim for objective measurements and eschew subjective or emotional dimensions. Calculating the probability of a future attack (2) also shares these sensitivities: in cybersecurity attackers are intelligent adversaries who may not follow patterns of previously seen behaviour. Consequently, calculating probabilities requires an element of subjective judgement, which raises questions when different opinions vary or conflict with objective data. There are a number of fallacies (3) that permeate cybersecurity [30], including the "sheer size fallacy" (*if one risk is smaller than an acceptable unrelated risk, then it should be accepted*), the "technocratic fallacy" (*since cybersecurity risks can be highly technical, only technical people can decide on what to do about them*), or the "fallacy of pricing" (*since we have to weigh the costs of risks against their benefits, it is necessary to place a monetary value on risks*). Finally, it is important to note that risks are not tolerated equally, and neither are they fairly and evenly distributed (4): those who make decisions about risks may not be impacted by them or pay for the costs associated with the risk; a company's underinvestment in security may result in significant harm to others (e.g. their customers) and not directly to themselves; one business's estimate of a tolerable level of risk may be deemed unacceptable by a regulator; a poorly implemented authentication solution may cause usability and productivity impacts to users, etc.

2.2.4 Principlist Frameworks. The last consist of principlist frameworks, an example of which was used in the highly influential Menlo report [2]. Principlism is a form of deontology, and principlist frameworks are articulated around a small number of fundamental principles derived from moral and professional ethical practices. These principles then drive what duties need to be satisfied, however complications arise when different duties are in conflict. A principlist framework thus aims to help navigate these issues by providing a lightweight means of helping to identify possible conflicts, however the exact nature, context, and importance of specific factors in these conflicts is left to the deliberation of researchers and practitioners. Furthermore, the resolution of such conflicts is also left open to the interpretation of the users of the framework. The Menlo report proposes four principles: Respect for Persons, Beneficence, Justice, Respect for Law and Public Interest. The first three of these are drawn from the Belmont report [56] that focuses on the protection of human research subjects, and the fourth is proposed as an additional category to highlight the wider legal and public interest in cybersecurity.

Formosa et al. [25] noted the growing importance of Artificial Intelligence and associated ethical issues [24] in the practice of cybersecurity. Drawing on the Menlo report and other principlist approaches [2, 37, 41, 57, 59], they propose a principlist framework

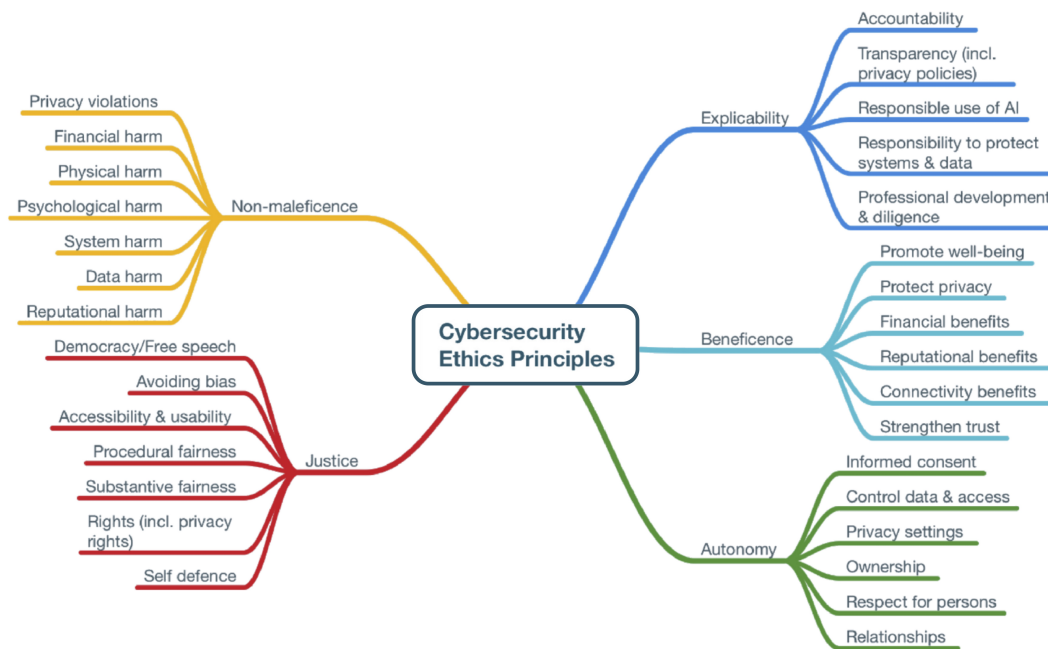


Figure 2: A principlist framework for cybersecurity ethics. Reprinted from [25] with permission from Elsevier

to address ethical issues according to five different principles, the definitions of which are reproduced here verbatim:

- *Beneficence: Cybersecurity technologies should be used to benefit humans, promote human well-being, and make our lives better overall.*
- *Non-maleficence: Cybersecurity technologies should not be used to intentionally harm humans or to make our lives worse overall.*
- *Autonomy: Cybersecurity technologies should be used in ways that respect human autonomy. Humans should be able to make informed decisions for themselves about how that technology is used in their lives.*
- *Justice: Cybersecurity technologies should be used to promote fairness, equality, and impartiality. It should not be used to unfairly discriminate, undermine solidarity, or prevent equal access.*
- *Explicability: Cybersecurity technologies should be used in ways that are intelligible, transparent, and comprehensible, and it should also be clear who is accountable and responsible for its use.*

This framework outlines different ethical values in cybersecurity, which can further be refined into more detailed concepts as illustrated in figure 2. Given that this is the first framework to explicitly include the consideration of the ethics of AI in cybersecurity, we chose this approach to help frame the ethical aspect of our investigation as described in section 3.

2.3 Cyber Security Body of Knowledge

Cybersecurity is a concept that has been defined and characterised in a variety of different ways, most of which frame it as a process for

protecting information by preventing, detecting, and responding to attacks [44]. Several properties of cybersecurity are regularly included in these definitions, the core three being confidentiality, integrity, and availability. Additional properties are sometimes added, such as authentication, non-repudiation, or utility, and the scope can also be wider than the protection of information to include computers, electronic communication systems, and electronic communications.

Going into more detail about the different areas of foundational and generally recognised knowledge that make up cybersecurity, the Cyber Security Body of Knowledge (CyBOK [46]) serves as a guide and maps the core elements of the discipline. At the time of writing, CyBOK has released v1.1 of its knowledgebase [18], which breaks down cybersecurity into five main categories: Human, Organisational & Regulatory Aspects; Attacks & Defences; Systems Security; Software and Platform Security; and Infrastructure Security. Within each of these categories, a total of 21 Knowledge Areas (KA) introduce and outline common material. Taken together, CyBOK can be used to understand the means and objectives of cybersecurity, mitigate against failures and incidents, and manage risks. Given that CyBOK provides a comprehensive breakdown of different areas of cybersecurity, we chose to use it to help frame our investigation as described in section 3.

2.4 Professional Codes of Practice

A number of professional bodies have adopted codes of practice to help their members navigate the ethical challenges that can arise in the performance of their duties. The Association for Computing Machinery (ACM) with approximately 100,000 members (educators, researchers, and professionals) is the world's largest computer society, and has published a detailed Code of Ethics and Professional

Conduct ([1]). This outlines and discusses the following general ethical principles:

- (1) Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing.
- (2) Avoid harm.
- (3) Be honest and trustworthy.
- (4) Be fair and take action not to discriminate.
- (5) Respect the work required to produce new ideas, inventions, creative works, and computing artefacts.
- (6) Respect privacy.
- (7) Honour confidentiality.

More specifically to cybersecurity, the Information Systems Security Association (ISSA) has also published a code of ethics [33], however this is quite short, consisting of the following:

- Perform all professional activities and duties in accordance with all applicable laws and the highest ethical principles;
- Promote generally accepted information security current best practices and standards;
- Maintain appropriate confidentiality of proprietary or otherwise sensitive information encountered in the course of professional activities;
- Discharge professional responsibilities with diligence and honesty;
- Refrain from any activities which might constitute a conflict of interest or otherwise damage the reputation of or is detrimental to employers, the information security profession, or the Association; and
- Not intentionally injure or impugn the professional reputation or practice of colleagues, clients, or employers.

The UK Cyber Security Council has its own code of ethics for cybersecurity grounded on the principles of integrity, professionalism, and credibility. In addition, the council has made available some guidance for individuals, articulated around Fair Competition, Honesty, Inclusion, Integrity, Lawful behaviour, Professionalism, Reporting, and Competence. Complementing this, the council published 16 case studies which illustrate the core ethical considerations. In the same way that we use CyBOK to frame our research into the ethical challenges of cybersecurity, the scenarios presented by the UK Cyber Security Council are also categorised according to the five main Knowledge Areas of CyBOK.

The Forum of Incident Response and Security Teams (FIRST) has also developed its own guidance in the form of a code of practice for professionals involved in incident response [23]. The code of practice outlines a number of duties (Duty of trustworthiness; Duty of coordinated vulnerability disclosure; Duty of confidentiality; Duty to acknowledge; Duty of authorization; Duty to inform; Duty to respect human rights; Duty to Team health; Duty to Team ability; Duty for responsible collection; Duty to recognize jurisdictional boundaries; Duty of evidence-based reasoning). In addition to providing a clear definition of these different duties, the code of practice also includes an appendix for dealing with dilemmas. This outlines how members may find themselves in a position where no action seems to satisfy all of the ethical principles. To address such dilemmas, practitioners are encouraged to reflect on how stakeholders may be affected by their actions and to favour solutions that minimize the infringement of the Code.

Overall, the guidance and codes of practice offered by professional bodies are helpful, but they typically fall short of the detail and nuance that comes from the ethical work outlined in Sections 2.1 & 2.2. Moreover, we note that beyond mapping exemplar ethical case studies to different areas of cybersecurity, it is not clear what kinds of ethical questions can arise in the exercise of specific professional duties. As a result, our research aims to investigate this by exploring the kinds of ethical challenges, insights, and solutions that arise in different areas of cybersecurity. In the following section, we describe the methodology we used to investigate this.

3 METHODOLOGY

Based on our review of the existing literature on the ethics of cybersecurity, we devised an interview guide to answer the research question: *“What ethical challenges, insights, and solutions arise in different areas of cybersecurity?”* We recruited 15 interviewees from outside our institutions through a mix of direct contact with existing connections and previous research participants who had consented to being contacted about future research. Our selection criteria required participants to have experience and knowledge of the cybersecurity profession.

We analysed the data using thematic analysis, which is an inductive coding process to help identify patterns in meaning from the data. Our analysis identified a number of themes and sub-themes pertaining to ethical concerns, experiences, and solutions that our participants related to different areas of cybersecurity.

The study was ethically reviewed and approved by the Departmental Research Ethics Committee at our institution.

3.1 Recruitment

We used several means to recruit our participants, including advertising on Twitter, Reddit, Mailing Lists and Blogs. We also reached out to participants on Slack channels and LinkedIn. To diversify our sample, we aimed to interview senior managers and executives who have likely made important security decisions. Since these are a hard-to-reach group [22], we used the snowball sampling method [27] to recruit some participants, and worked with a consultant advisor who had wider access to senior executives working in security. We note that the results of the convenience sampling cannot be generalised to the target population because of the potential bias of the sampling technique [3] and cannot be used to identify differences of population subgroups [3].

All of our participants were working at different companies. At the time of recruitment, interested participants were employees who were active at their company.

We asked interested participants to complete an online screening questionnaire. We received 85 complete responses. In addition to asking demographic questions, we asked participants to provide details on their employment as well as their company size.

We selected participants based on those who could best inform our research question and enhance our understanding of the ethical challenges, insights, and solutions arising in cyber security. Hence, we chose participants based on occupation (e.g. technical, managerial), field (e.g. malware, privacy, web security), relevance to CyBOK's knowledge base (e.g. human, organisational, and regulatory aspects), experience level (e.g. junior, senior), and diversity

(e.g. gender, ethnicity). We note that all our participants worked in the private sector and none worked directly for the government or government affiliated employers.

We describe the demographics of our participants in Table 1.

3.2 Pilot Study

To validate our initial interview questions, we conducted a pilot study with two individuals in our research institution. We recruited the pilot participants through snowball sampling. Two researchers analysed the pilot interviews. We used the findings to identify potential problems (e.g., adverse events, time) in advance prior to conducting the full-scale study. We didn't use the results from the pilot interviews, but we have refined our interview questions to ensure they are non-leading and clearer for participants.

3.3 Demographics

Table 1 summarises the demographics of our sample (n=15). We interviewed nine male and five female participants. Ages ranged from 18 to 55. Eight participants were interviewed in person, and seven remotely.

Additionally, we collected more contextual information about our participants and the business sector in which they operate (socioeconomic status of customers, types of products offered, geographic locations served). We conducted 5 interviews in-person (in secure locations in our institution) and 10 interviews remotely (on Microsoft Teams).

Table 1: Study Demographics

P#	Age (M/F)	Occupation	Cyber Security Field	CyBOK
P01	50-55 (M)	Managing Director	Privacy & Online Rights	[55]
P02	45-50 (M)	Systems Security Lead	Secure Software Lifecycle	[60]
P03	45-50 (F)	Security Manager	Risk Management & Governance	[5]
P04	35-40 (M)	Security Engineer	Security Operations & Incident Management	[19]
P05	40-45 (M)	Cloud Security Lead	Network Security	[49]
P06	40-45 (F)	Infosec Lead	Web & Mobile Security	[21]
P07	30-35 (M)	Security Engineer	Malware & Attack Technologies	[36]
P08	30-35 (F)	Freelance Ethical Hacker	Adversarial Behaviours	[54]
P09	45-50 (M)	Product Security Lead	Forensics & Software Security	[45, 50]
P10	18-25 (M)	Security Developer	Malware & Attack Technologies	[36]
P11	30-35 (M)	Security Engineer	OS & Virtualisation Security	[4]
P12	25-30 (F)	Penetration Tester	Web & Mobile Security	[21]
P13	35-40 (F)	Security Manager	Authentication & Authorisation	[26]
P14	30-35 (M)	Security Consultant	Physical Layer & Telecom Security	[58]
P15	25-30 (M)	Privacy Engineer	Human Factors & Law & Regulation	[8, 51]

3.4 Procedure

3.4.1 Semi-structured Interviews. We followed a semi-structured interview protocol utilising an interview guide to maintain direction while keeping the interview open for both depth and breadth of topic exploration. In order to help focus our questions on specific aspects of cybersecurity professions, we made use of CyBOK's five broad categories:

- (1) Human, Organisational & Regulatory Aspects (e.g. risk management & governance, law & regulation, human factors, privacy & online rights)
- (2) Attacks & Defences (e.g. malware & attack technologies, adversarial behaviours, security operations & incident management, forensics)
- (3) Systems Security (e.g. cryptography, operating systems & virtualisation security, distributed systems security, formal methods for security, authentication, authorization & accountability)
- (4) Software and Platform Security (e.g. software security, web & mobile security, secure software lifecycle)
- (5) Infrastructure Security (e.g. applied cryptography, network security, hardware security, cyber physical systems, physical layer & telecommunications security)

In addition, to help focus on different aspects of cybersecurity ethics, we used the five principles proposed by [25], outlined above in Section 2.2. Our interview guide is included in Appendix A.

3.4.2 Thematic Analysis. The interview data was analysed using Thematic Analysis. According to [16], it is a common method of analysis in qualitative research and involves identifying, analysing, interpreting, and reporting patterns of meaning (known as themes or codes) from qualitative data. Thematic analysis is also frequently used with existing theoretical frameworks to provide interpretive power. Given that our approach made use of the principlist framework for cybersecurity ethics proposed by [25], thematic analysis proved highly suitable in this regard, and helped to provide greater insight into detailed ethical considerations.

Two researchers were involved in the data collection and analysis. The primary researcher, who conducted most of the interviews, did an initial coding of the interview transcripts. To ensure credibility of the codes, a second researcher cross-checked all the codes against the interview transcripts. Any differences and/or issues arising from the coding were discussed and resolved among the two researchers. A codebook consisting of 122 codes emerged from the initial coding. These codes were then applied across other interviews through constant comparison, while new codes were added as they emerged and were deemed necessary. In further analysis, the researchers discussed and grouped the codes into themes. Regular coding meetings were held to discuss any emerging codes and to group the codes into families.

We observed data saturation [17, 28, 52] between the 13th and the 15th interview; i.e. no additional issues or insights emerged from data and all relevant conceptual categories had been identified, explored, and exhausted. Hence, we stopped interviewing.

3.5 Research Ethics

The University of Oxford's Central University Research Ethics Committee (CUREC) reviewed and approved the study (C1B-23HT-COML-003). Prior to each interview, participants were briefed and signed an informed consent form explaining our study and data confidentiality practices. Due to the sensitivity of our interviews, we asked participants not to name specific people or sites so that the interviews will be anonymous to some degree.

All interviews were AES 256 encrypted and stored in a physical safe in our organisation. Participants were thanked for their time with GBP £50 in electronic store vouchers. In addition, participants were reimbursed for out-of-pocket expenses related to participation, including travel, meals, accommodation, and childcare. Participants could withdraw themselves and their data at any point, without loss of compensation, and without providing a reason. No participant withdrew.

4 FINDINGS

In this section, we detail the findings of our study. We discuss our key findings organised according to the main themes of our analysis. The main themes are:

- Human, Organisational, Regulatory Aspects (§4.1)
- Attacks & Defences (§4.2)
- Systems Security (§4.3)
- Software & Platform Security (§4.4)
- Infrastructure Security (§4.5)

4.1 Human, Organisational & Regulatory Aspects

Our participants reported a number of ethical concerns, dilemmas and challenges in relation to the Human, Organisational and Regulatory Aspects of cyber security which encompasses Risk Management and Governance, Law and Regulation, Human Factors and Privacy and Online Rights. Participants experience ethical challenges related to the constant need for maintaining confidentiality (§4.1.1), balancing the competing interests of protecting their company's reputation and maintaining user security (§4.1.2), and disclosing security risks without making users feel insecure (§4.1.3).

4.1.1 Always Maintaining Confidentiality. Our participants (n=5) stated that due to the nature of their profession, they handle and see private, sensitive, and proprietary information that must be kept entirely secret – adding that maintaining confidentiality is highly critical. They are instructed to maintain the confidentiality of information they come upon, or face significant ramifications for their career. However, they experienced ethical dilemmas and challenges deciding between maintaining or breaking confidentiality. Participants struggled to carefully weigh the ethical implications and informed decisions of whether or not to disclose confidential information. Our participants stated this is a complex challenge that is dependent on many factors such as whether there is knowledge about harm, crime or illegal activities, and whether the confidential information should be disclosed to law enforcement or other government agencies. For instance, participant P07, who works at a cybersecurity and anti-virus company gave more insights on the

ethical challenges to maintaining confidentiality while repairing their client's machines:

"We are asked to fix people's computers all the times. Our job is getting rid of the virus, not looking at our clients' photo albums or tax returns. We don't care what data is on their computer, we're only here to repair it. I'm not saying we do, but if we ever see any confidential information, I assume it can get quite complicated. If someone is cheating on their spouse or avoiding taxes, then it's none of our business. But if we find by accident that there is threat to human life or a child in danger, we have an obligation to report it." – P07, Security Engineer

4.1.2 Conflicts between Business & Security Practices. Our participants (n=7) reported that conflicts of interest frequently arise in cyber security firms due to the competing interests between business and cyber security practices. They reported that conflicts of interest emerge between individual interests, public interests and corporate interests. Some of our participants suggested that cyber security practices should be prioritised over business or profit-making activities, as they strengthen the company's reputation and lead to trust over the long term. A common topic of discussion was the ethical challenges of dealing with data breaches: some companies tend to be reluctant to inform the public of data breaches in an attempt to protect their brand. For instance, Participant P01, who manages a company that specialises in data privacy management tools, provided more insights on the conflicts of interests arising between data breaches and reputational damage:

"We had companies where they've had multiple breaches, seriously embarrassing breaches. The company is still trading and still has actually gone from strength to strength. When they had the breach, all they did was lower their prices for a while and then bring it up again. They actually increased their market reach. CEOs will then tell you: "We got a plan to deal with it if it occurs". They don't really care about the breach. They obviously care about the reputational damage. They've got a plan on how to deal with the reputational damage. They're actually not addressing the issue." – P01, Managing Director

Moreover, Participant P03 who is manager at a small cyber security firm stated that a key ethical challenge is whether to conduct business with clients who are not willing to invest in adequate cyber security solutions.

4.1.3 Disclosing Security Risks Without Making Users Feel Insecure. Our participants (n=4) reported ethical challenges, and dilemmas in disclosing security risks without making their customers or users feel insecure. Our participants reported that this can be a challenging balancing act. On one hand, disclosing security risks is a transparent business practice that improves transparency and ensures that users take informed decisions. On the other hand, disclosing security vulnerabilities can risk making users feel unnecessarily insecure or anxious about their security. For instance, Privacy Engineer P15 provided more details on this challenge:

“Sometimes you don’t want to put the end-user in a state where they’re not so sure about the systems and how their data is being processed and stuff just because there’s a minor issue which you are able to pick on your own. So it’s very hard sometimes when you’re looking at the business side as well as the ethical side.” – P15, Privacy Engineer

4.2 Attacks & Defences

Our participants reported a number of ethical concerns, dilemmas and challenges in relation to the Attack & Defence aspects of cyber security which encompasses Malware & Attack Technologies, Adversarial Behaviours, Security Operations & Incident Management and Forensics. Participants experience ethical challenges related to ethical security hacking and cyber intrusion (§4.2.1), and defending against cyber attacks (§4.2.2).

4.2.1 Ethical Security Hacking and Cyber Intrusion. Our participants (n=3) reported concerns and dilemmas in relation to ethical security hacking and cyber intrusion. Some of our participants were ethical hackers hired by different companies to test their security systems and find security vulnerabilities. They were responsible for applying offensive cyber tools and techniques to identify and drive security improvements. In return, they would reveal vulnerabilities to their client and create guidelines on how to address them, thus helping to secure company networks to protect trade secrets and business practices. Our participants reported that their most ethical concerns or challenges result from the possibility of finding illegal activity or unethical practices connected to their client. For instance, participant P08, a freelance ethical hacker who delivers penetration testing and red teaming capabilities for companies, discussed the ethical ramifications of the possibility of discovering illegal activity during a penetration test.

“During a pentest, you could come across illegal activity. This is where it gets messy. My clients tend to be from all over the world, what is legal in the UK might be illegal in other countries. Do you report it to the authorities? You could get investigated yourself. If you report it, are you breaking your NDA? That all depends on the jurisdiction. Do you not report it? You could be complicit to a crime you don’t report. I am personally aware of pentesting companies that have a provision allowing them to report illegal activities observed.” – P08, Freelance Ethical Hacker

Moreover, participant P12, who is a penetration tester, stated that they evaluate their own values and moral principles; and don’t take any projects that could create a conflict with their own or with other societal ethical values.

While this finding is similar to the one reported in Section 4.1.1, we note that it is more grounded in the context of ethical hacking through a client-freelancer relationship, whereas the finding in Section 4.1.1 relates more to ethical dilemmas faced by employees working in larger companies.

4.2.2 Defending Against Cyber Attacks. Our participants (n=2) reported ethical concerns and challenges in relation to defending

against cyber attacks. Some of our participants were affected by remote attacks (e.g. DDoS) and were concerned about the use of force to defend against cyber attacks. In some cases, our participants reported the need urgently to take down malicious servers that were conducting attacks on their infrastructure. While our participants had an ethical obligation to protect their organisation’s infrastructure, they had significant ethical concerns over the use of force to defend against cyber attacks. Participants were concerned that the use of force may create unacceptable risk to violate the rights of innocent individuals and organisations, since the collateral effects of the use of force are unknown. For instance, participant P04, who works as a security engineer at a large company, explained how they dealt with DDoS against their company’s servers:

“We get DDoS attacks all the time, these get automatically blocked with our firewall. If the attack is overwhelming, we’d never retaliate. We go through legal means. Sometimes this means sending abuse reports to the host company. Quite recently, we were dealing with a bunch of IP addresses that were being used to facilitate DDoS amplification attacks. We reported them to the host company and they took them down very quickly.” – P04, Security Engineer

Moreover, Security Developer P10 stated that they avoid the use of force to defend against cyberattacks because cyberattacks often involve ‘spoofing’ strategies which make it easy to misidentify the system responsible for the attack.

4.3 Systems Security

Our participants reported a number of ethical concerns, dilemmas and challenges in relation to the Systems Security aspects of cyber security which encompasses Cryptography, Operating Systems & Virtualisation Security, Distributed Systems Security, Formal Methods for Security and Authentication, Authorisation & Accountability. Participants experience ethical challenges related to cybersecurity resource allocation (§4.3.1), and contracting third party security services (§4.3.2).

4.3.1 Cybersecurity Resource Allocation. Our participants (n=4) reported that there was an ethical concern around the tradeoff between security and other functionalities or priorities. They explained that some cybersecurity solutions (i) may consume considerable individual and organisational resources such as time, labour, money, and expertise; (ii) may negatively impact data storage capacities, bandwidth (upload/download) speeds, energy usage, and the usability and reliability of systems. As such, they experienced challenges in making ethical decisions about how to allocate cyber security resources to protect the security of their customers without creating significant burdens (cost, convenience, and functionality). On one side, not having effective cyber security solutions might cause serious harm and damage to users’ security. On the other side, extreme cyber security solutions might be unusable or economically unsustainable. For instance, security manager P13 discussed the challenges of balancing password security policies:

“Balancing this is never easy. We made employees change their passwords every 90 days after one employee had

their passwords leaked. The point was to improve our security posture, but that didn't work out. Employees were writing their passwords everywhere." – P13, Security Manager

4.3.2 Contracting Third Party Security Services. Our participants (n=9) reported that they had ethical concerns about the use of third party security services. Contracting third party security services (for purposes such as security auditing, penetration testing, security) was effective and cost-efficient for our participants. However, they were concerned about the ethical behaviour of third party services, especially those who operate from countries with challenging political and legal climates. Our participants reported that they exercised due diligence to ensure that contracted third parties are responsible and committed to high standards of ethical conduct. However, some participants reported ethical obstacles that occurred from contracting third parties security services. For instance, Product Security Lead P09 who contracted a remote penetration testing service to evaluate the security of a computer-forensic online service reported a potential inappropriate ethical behaviour.

"We hired this pentesting firm and signed all the paperwork which clearly said that if they find any vulnerabilities, they need to tell us immediately. We discovered later that they had found but failed to report major vulnerabilities in our website. We contacted them multiple times regarding this, and they never got back to us." – Product Security Lead, P09

4.4 Software and Platform Security

Our participants reported a number of ethical concerns, dilemmas and challenges in relation to the Software and Platform Security aspects of cyber security which encompasses Software Security, Web & Mobile Security, Secure Software Lifecycle. Participants experience ethical challenges related to disclosure and patching of vulnerabilities (§4.4.1), and prioritising vulnerability patching practices (§4.4.3).

4.4.1 Disclosure and Patching of Vulnerabilities. Our participants (n=3) reported ethical concerns and tensions between disclosure and patching of vulnerabilities. Our participants had an ethical duty to be transparent about vulnerabilities found on their system – so that affected parties can make informed decisions. However, they were wary that disclosing vulnerabilities could make it easier for malicious or bad actors to exploit them. As such, they experienced ethical challenges in balancing between the need for security and the need for transparency. For instance, Systems Security Lead P02 stated that their organisation publicly discloses details of vulnerabilities in their products. However, this is done at the discretion of the organisation and based on a coherent ethical judgement about what is best to do, given the facts, options, products and interests at stake. They explained:

"There is no one-size-fits-all approach. You have to weigh the risks and benefits involved. A common practice is to disclose them after they have been patched. But, if a vulnerability is critical and can cause severe harm, you might need to disclose it as soon as possible." – Systems Security Lead, P02

4.4.2 Prioritising Vulnerability Patching Practices. Our participants (n=6) reported that a common ethical challenge arises when prioritising vulnerability patching practices. Our participants reported that they have an ethical duty to respond to discovered vulnerabilities in a timely manner. Due to the high number of vulnerabilities reported, our participants have to prioritise which vulnerabilities and assets to patch and in what order. Our participants revealed that they often interpret vulnerability risk metrics subjectively and raise a possible dilemma when weighing whether to protect the organisation's interests over those of customers. For instance, Information Security Lead P06 revealed that it is ethically challenging to prioritise between patching vulnerabilities that have been actively exploited in the wild and patching vulnerabilities that are affecting company assets. They said:

"It is always challenging to choose which vulnerabilities to patch first because they always lack context. But even if you got more context, you have to make difficult decisions. Do you patch a vulnerability that is being exploited and affecting your users' data first? Or do you patch a vulnerability that can expose your employees' details?" – P06, Information Security Lead

4.4.3 Disclosing Security Incidents Without Losing Customer Trust. Our participants (n=3) reported ethical challenges, and dilemmas in disclosing security incidents that have already been addressed and no longer pose a security threat to users and customers. Our participants expressed concerns that disclosing resolved incidents may damage customer trust, and could result in reputational damage to the company or the relevant security team. As such, participants reported carefully considering the ethical implications of disclosing already-solved incidents, and made informed decisions. Patched vulnerabilities were often disclosed but were kept confidential in some cases. Privacy Engineer P15 explained:

"At that point in time you're like should we or should we not tell the end users? Sometimes there's an incident which is not visible to the end user and you are able to detect and fix it internally on your end. And just deciding as in: Should we tell the users that this happened or should we just remain silent? So you're trying to weigh the cost, you're doing the cost-benefit analysis between telling the end-user the impact of the business that's going to bring. So sometimes you can end up just like: Hey, fix it internally but don't even tell the end user, just keep quiet." – P15, Privacy Engineer

4.5 Infrastructure Security

Our participants reported a number of ethical concerns, dilemmas and challenges in relation to the Infrastructure Security aspects of cyber security which encompasses Applied Cryptography, Network Security, Hardware Security, Cyber Physical Systems and Physical Layer and Telecommunications Security. Participants experienced ethical challenges related to balancing infrastructure security with privacy (§4.5.1), and ensuring accountability and responsibility in AI (§4.5.2).

4.5.1 Balancing Infrastructure Security with Privacy. Our participants (n=2) reported that they faced dilemmas in balancing the

security of infrastructure systems with the need to maintain the data protection and privacy of their users. Infrastructure security systems often collect and store data for a variety of purposes, such as protecting users, identifying potential threats, tracking system performance, and improving security measures. However, participants were wary of the privacy and ethical ramifications of using the data of the customers. Security Consultant P14 who works as a security consultant for a company that manages internet infrastructure discussed how they take steps to protect the security of their customers without compromising their privacy. They explained:

“We could scan customer data but we don’t want to do that because it involves essentially snooping on customers traffic. On the flip side of that however, if you move the perspective away from the customer’s edge access in their homes, the edge of the network, what we will do is identify and be provided with a list from various agencies and commercial and government groups of known bad URLs and IP addresses, and we can actually drop the traffic in the core network. If you imagine our routing and switching network in London, for example, if we see packets come in on our consumer network, we don’t care where they’ve come from, we don’t want to know who sent them into the network, we’ll see the IP address that we know is bad or the destination is bad, and we drop the packets there. The consumer is protected because they try to go to a website or service, which is malicious, but we didn’t do the content filtering at the edge on their premises.” – Security Consultant, P14

4.5.2 Ensuring Accountability and Responsibility in AI. Our participants (n=9) reported a number of security concerns in infrastructure security related to ensuring accountability and responsibility for the use of AI in infrastructure security. They explained that existing infrastructure can often be complex, opaque and has the potential for abuse of power and lack of accountability. Moreover, participants indicated that some infrastructure security systems rely on AI surveillance to identify and track potential threats, which has the potential for mass surveillance and violating user privacy without accountability. Other participants raised ethical concerns about potential for bias in AI systems and the lack of transparency in how AI systems make decisions. Cloud Security Lead P05 explained:

“Many companies are now using AI in security technologies and then selling these to governments who then integrate them in critical infrastructure. I don’t know if people are aware that AI systems can be attacked. And these attacks can have significant severe effects. I think that is one of the biggest ethical concerns in cybersecurity.” – P05, Cloud Security Lead

We note that unlike other findings, this is a primarily security concern related to the exploitation of bias and discrimination in artificial intelligence.

5 DISCUSSION

5.1 Ethics of Cybersecurity Decision-Making

To help frame our discussion, we propose that cybersecurity can be defined as *desirable decisions and actions to manage threats, vulnerabilities, and impacts*. While cybersecurity exists within a legal and regulatory framework, this definition helps us to focus on two fundamental concepts that are central to the ethical challenges that cybersecurity professionals can face: what is *desirable* and how are *decisions* made.

Desirable relates to perception and judgement of what is and what is not wanted but, crucially, cybersecurity always has multiple different stakeholders who may have different perceptions and priorities on matters of security and how to control potential problems. This can clearly lead to conflicts that require resolution in order to act, however it is important to remember that cybersecurity has subjective, objective, and affective dimensions that are specific to each relevant stakeholder. Thus, being able to understand the varied nuances and different perspectives between the objective views on security, the subjective interests of multiple stakeholders, and the emotional implications of cybersecurity actions is necessary. We also note that cybersecurity is a highly technical subject and that cybersecurity actions entail the exercise of power and control over others in various contexts. As a result, different stakeholders have varying levels of security understanding, and also hold cultural, social, and normative views which influence what they perceive to be threats, and what are desirable ways of managing this.

Moreover, in cybersecurity the authority to decide on actions is vested in an individual or small group of individuals who usually (but not always) have the accompanying responsibility of ensuring the success of cybersecurity activities. In order to achieve optimum outcomes, these decision-makers need to resolve complex problems that require a balance between the technical aspects of cybersecurity (threat, vulnerability, impact) and the perceptions and interests of multiple different stakeholders (including subjective and affective dimensions), which is guided by overarching principles. Many of these principles are grounded in “best practice” in cybersecurity (e.g. the principle of least privilege, defence in depth, separation of duty, etc.), however these principles also encompass ethical values (e.g. beneficence, justice, non-maleficence, explicability, or autonomy as noted in Section 2).

Any decision in cybersecurity is therefore driven by the ability of a decision-maker to reach a balance between their own and multiple other stakeholder perspectives on how to manage threats, vulnerabilities and impact, which covers areas such as compliance with relevant laws and regulation, being sensitive to contextual issues (such as cultural, social or religious norms), and adhering to fundamental security, privacy and ethical principles. While practitioners can rely on precedent and “best practice” to navigate these decisions, it is readily apparent that there are a number of issues that are not covered by this. Specifically, while precedent can provide helpful framing and information pertaining to past decisions, cybersecurity is an evolving problem: new threats and vulnerabilities get discovered, new security technologies and practices are devised, the likelihood of attacks happening can change over time, and so can contextual factors such as attitudes, practices and expectations. Moreover, while there is a strong desire for “best practice” or

“tried-and-tested” security, this typically results in a “solutioneering” approach: where a pre-existing cybersecurity solution is applied to multiple different problems and contexts without the understanding required for making appropriate decisions. Furthermore, the insistence on “tried-and-tested” security means that innovation is discouraged and improvements are slow to be recognised and adopted.

5.2 Decisions in Cybersecurity Professions

As our findings indicate, decisions in cybersecurity can be made by many different individuals across the profession. Using CyBOK, our findings help map out ethical issues that arise in different areas of the cybersecurity profession, and provides a useful starting point for a more systematic exploration. The issues identified in section 4 highlight several issues that arise when making specific decisions: difficulties in balancing different perspectives on prioritising vulnerability patching, navigating the emotional aspects of customer trust that can be impacted by disclosing information, or resolving conflicting views about the contrasting interests of the business against the interests of the customers. More research is needed to systematically map the different types of decisions that arise in the pursuit of professional duties, and help to provide a more specific breakdown of the kinds of ethical dilemmas that may arise in each relevant area of the profession. This would be beneficial in cataloguing the variety of security decisions that pose ethical dilemmas, informing curriculum design and educational material for cybersecurity professionals, and providing deeper opportunities to examine cybersecurity ethical issues in more detail. While the ethics of cryptographic research [48] or vulnerability disclosure [40] have been highlighted and investigated, we identified a number of other ethical challenges that need greater scrutiny, such as considering the ethical implications of associating with irresponsible third parties (either clients or service providers), possible harm arising from how security information is interpreted by customers, or navigating conflicts between company and customer interests.

One dimension that was not identified explicitly by our participants is the inherent potential for them to be placed in a personal conflict of interest when making such decisions: for example when a possible decision outcome would result in a negative personal impact but a much larger benefit for others, or when weighing the convenience of simply reusing a pre-existing solution compared to a more detailed evaluation of how suitable it is to the context of use.

5.3 Ethical Cybersecurity AI

Finally, it is important to recognise the growing concern over the use and applicability of AI technology in cybersecurity. It is foreseeable that this technology will be used in ways to improve the quality, timeliness and relevance of information available to decision-makers and also to help them identify appropriate solutions. It is also likely that AI models will be used and increasingly be relied upon to make decisions about detailed technical or time-critical cybersecurity problems, however a fundamental concern about AI models pertains to their lack of transparency and explicability. AI decision tools will need to account for the complexity of different perspectives, values, principles outlined above, however it is likely

that they too will be subject to bias, unfair, and even malicious influence. Moreover, given the complexity and lack of transparency inherent in these approaches, such biases, malicious influences, or conflicts of interest may become harder to identify, challenge, and remediate.

5.4 Implications for Professional Practice

As noted in Section 2.4, the current codes of practice from professional bodies do not provide very detailed or nuanced guidance to help practitioners navigate ethical dilemmas in the pursuit of their duties. As discussed above, we believe more research is needed to map and explore the ethical challenges that arise in different areas of the cybersecurity profession, however it is challenging to make recommendations for how these problems should be addressed on a more pragmatic level. Addressing the need for greater consideration of ethical issues through legislation is unlikely to succeed due to the complex and varied nature of moral dilemmas which can arise in cybersecurity. However, insights can be drawn from the long history of medical ethics and practice: from the early origins of the Hippocratic Oath to more modern Codes of Medical Ethics. In particular, we note that due to the manner in which the medical profession is regulated (e.g. through the General Medical Council in the UK) violations of such codes of conduct can lead to disciplinary proceedings and result in the loss of a license to practice medicine. As cybersecurity codes of ethics continue to be developed and updated, we think that it is also important to consider and investigate how ethical breaches could lead to disciplinary proceedings and affect professionals' license to practice cybersecurity.

Furthermore, we believe it is crucial for ethical awareness and training to be integrated into and throughout the education of cybersecurity professionals. This should complement work towards embedding ethical and responsible values into the principles, practices and frameworks that collectively make up the cybersecurity profession. Both of these recommendations are highly consistent with the aims of CyBOK, and we argue that (1) more work needs to be done to identify, map, and inform the ethical implications of cybersecurity practice in each of the CyBOK Knowledge Areas; (2) greater efforts need to be made to influence the curriculum, education, and training of cybersecurity professionals to cover and help them manage to navigate ethical dilemmas; (3) more should be done to share and learn from how ethical dilemmas have been tackled by practitioners, including making room for more “honest” and “safe” spaces for exchange among security officers in relation to the ethical dilemmas they encounter. Finally, (4) it is imperative for more research to be undertaken into the ethics of cybersecurity AI applications, in order to identify and frame the possible benefits, disadvantages, and dangers of this technology.

6 LIMITATIONS

Our study has some limitations. First, when we engaged with our participants, we devised an interview guide that was informed by two different existing pieces of work: CyBOK and the principlist framework proposed by Formosa et al. [25]. Despite our best intentions in explaining and grounding our engagement in ethical principles, it is evident that much of this detail was overlooked in favour of discussing example cases of ethical challenges. Upon

further reflection, it is clear that a principlist approach to empirical ethical research presents challenges in how to engage with participants, however we believe that such principles are very helpful in interpreting the examples and comments given by our participants. In contrast, the detailed breakdown of different areas of cybersecurity from CyBOK provided a much more helpful framing for our participants and proved a useful means of prompting for different experiences of ethical issues across the whole spectrum of cybersecurity practice.

Second, common to all qualitative studies, researcher bias is a concern. Both researchers were trained to conduct research interviews, taking care to avoid leading questions, and ensuring that participants felt comfortable to respond to questions. The researchers avoided interrupting participants, and probed for more information when required.

Third, given that cybersecurity and ethics are sensitive topics, participants may have been concerned about sharing information, or otherwise been biased in their responses. In order to allay these concerns, an information sheet was given to all participants which described the details of the research and how their data would be used. Each participant was verbally invited to ask questions about this sheet prior to the interviews starting, and we emphasised that their comments would be anonymised and the interview data encrypted and protected. While it is difficult to avoid biased answers when undertaking interview research, we tried to account for this in our analysis of the data by carefully reviewing the codes and themes we identified. Given that our research was positioned as a high level exploration of ethical issues in different areas of the cybersecurity profession, we did not feel it necessary to take additional steps to avoid biases at this stage.

Fourth, the number of interviews is relatively small, and there may be concerns about the validity or generalisability of these findings more widely. While qualitative studies do not aim for statistical representativeness, they do aim to identify conceptual tendencies that can be considered more widely. In keeping with other forms of qualitative research, our aim here is to gain an in-depth understanding of the participants' views on ethics according to different areas of the cybersecurity profession. It is from this detailed understanding that our themes have emerged, and we argue that these can form the basis of future research aiming to generalise this to wider populations and the broader field of cybersecurity ethics.

Finally, in common with other work in empirical ethics, our qualitative findings are limited to the experiences and views of our participants. In particular, given that many of our participants had years of experience working in industry, it is not likely that they would have benefited from more recent efforts at teaching ethics as part of the curriculum of cybersecurity education. Furthermore, our discussion necessarily incorporates our normative views – while we have endeavoured to ensure that relevant information from our data and appropriate prior work and relevant theory has been presented, the choice and conclusions we draw are strongly influenced by our views as researchers.

7 CONCLUSION

In this paper, we have presented a review of existing work in cybersecurity ethics, and the results of our investigation into “What ethical challenges, insights, and solutions arise in different areas of cybersecurity?” We interviewed 15 cybersecurity professionals to investigate this, and found a number of different challenges which we analysed and mapped to the different areas of the Cybersecurity Body of Knowledge. Our findings highlight the inherent complexity and nuance associated with making decisions in cybersecurity, and we discussed how different objective, subjective and affective perspectives need to be taken into account, together with balancing technical, security, and ethical considerations. We further outlined the challenges for professionals who make such decisions and concluded that future research should investigate how ethical values can be embedded throughout the cybersecurity profession, including its tools, practices, and processes. In addition, it is imperative for ethical training and education to become part of the cybersecurity curriculum in order to provide professional practitioners with the means of identifying and addressing the ethical challenges that will arise in the performance of their duties, and we believe that CyBOK offers a very helpful framework for doing so.

Our participants also reported widespread concerns associated with the use of AI in cybersecurity, and we discussed how the technology offers opportunities to support ethical decision-making by providing greater insights and information into the problems being addressed. We also identified that the challenges of making complex cybersecurity decisions are likely to drive the application and adoption of AI-based decision-making. Given the complexity, subjectivity, and contextual nature of cybersecurity decision-making, and coupled with concerns around lack of transparency, possible biases, and potential for being maliciously influenced, we urge the research community to continue and expand the work undertaken in exploring the ethics of cybersecurity AI.

8 ACKNOWLEDGEMENTS

This work was supported by the Cyber Security Body of Knowledge (CyBOK) call for funded projects to develop resources around CyBOK v1.1 ©Crown Copyright 2023

REFERENCES

- [1] ACM. 2023. ACM Code of Ethics and Professional Conduct. <https://www.acm.org/code-of-ethics>
- [2] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The menlo report. *IEEE Security & Privacy* 10, 2 (2012), 71–75.
- [3] Marc H. Bornstein, Justin Jager, and Diane L. Putnick. 2013. Sampling in developmental science: Situations, shortcomings, solutions, and standards. *Developmental Review* 33, 4 (2013), 357–370. <https://doi.org/10.1016/j.dr.2013.08.003>
- [4] Herbert Bos. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Operating Systems & Virtualisation. <https://www.cybok.org/KA Version 1.0.1>.
- [5] Pete Burnap. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Risk Management & Governance. <https://www.cybok.org/KA Version 1.1.1>.
- [6] Sam Burnett and Nick Feamster. 2015. Encore: Lightweight measurement of web censorship with cross-origin requests. In *Proceedings of the 2015 ACM conference on special interest group on data communication*. 653–667.
- [7] John W Byers. 2015. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests—Public Review. In *Technical Report*.
- [8] Robert Carolina. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Law & Regulation. <https://www.cybok.org/KA Version 1.0.2>.

- [9] George Chalhoub. 2022. *The UX of Things: Exploring UX Principles to Inform Security and Privacy Design in the Smart Home*. PhD Thesis. University of Oxford. Available at https://www.georgechalhoub.com/pdf/chalhoub_thesis_2022.pdf.
- [10] George Chalhoub and Ivan Flechais. 2020. "Alexa, Are You Spying on Me?": Exploring the Effect of User Experience on the Security and Privacy of Smart Speaker Users. In *HCI for Cybersecurity, Privacy and Trust: Second International Conference, HCI-CPT 2020, Held as Part of the 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings* (Copenhagen, Denmark). Springer-Verlag, Berlin, Heidelberg, 305–325. https://doi.org/10.1007/978-3-030-50309-3_21
- [11] George Chalhoub and Ivan Flechais. 2022. Data Protection at a Discount: Investigating the UX of Data Protection from User, Designer, and Business Leader Perspectives. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 436 (nov 2022), 36 pages. <https://doi.org/10.1145/3555537>
- [12] George Chalhoub, Ivan Flechais, Norbert Nthala, and Ruba Abu-Salma. 2020. Innovation Inaction or in Action? The Role of User Experience in the Security and Privacy Design of Smart Home Cameras. In *Proceedings of the Sixteenth USENIX Conference on Usable Privacy and Security (SOUPS'20)*. USENIX Association, USA, Article 11, 20 pages.
- [13] George Chalhoub, Ivan Flechais, Norbert Nthala, Ruba Abu-Salma, and Elie Tom. 2020. Factoring User Experience into the Security and Privacy Design of Smart Home Devices: A Case Study. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3334480.3382850>
- [14] Markus Christen, Bert Gordijn, and Michele Loi. 2020. *The ethics of cybersecurity*. Springer Nature.
- [15] Markus Christen, Bert Gordijn, Karsten Weber, Ibo van de Poel, and Emad Yaghmaei. 2017. A Review of Value-Conflicts in Cybersecurity: An assessment based on quantitative and qualitative literature analysis. *The ORBIT Journal* 1, 1 (2017), 1–19. <https://doi.org/10.29297/orbit.v1i1.28>
- [16] Harris Ed Cooper, Paul M Camic, Debra L Long, AT Panter, David Ed Rindskopf, and Kenneth J Sher. 2012. *APA handbook of research methods in psychology, Vol 2: Research designs: Quantitative, qualitative, neuropsychological, and biological*. American Psychological Association.
- [17] Juliet Corbin and Anselm Strauss. 2014. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Sage publications.
- [18] CyBOK. 2023. Cyber Security Body of Knowledge. <https://www.cybok.org/>
- [19] Hervé Debar. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Security Operations & Incident Management. <https://www.cybok.org/ KA Version 1.0.2>
- [20] Anirudh Ekambaranathan, Jun Zhao, and George Chalhoub. 2023. Navigating the Data Avalanche: Towards Supporting Developers in Developing Privacy-Friendly Children's Apps. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 2, Article 53 (jun 2023), 24 pages. <https://doi.org/10.1145/3596267>
- [21] Sascha Fahl. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Web & Mobile Security. <https://www.cybok.org/ KA Version 1.0.1>
- [22] Jean Faugier and Mary Sargeant. 1997. Sampling hard to reach populations. *Journal of Advanced Nursing* 26, 4 (1997), 790–797. <https://doi.org/10.1046/j.1365-2648.1997.00371.x>; eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1365-2648.1997.00371.x>
- [23] FIRST Ethics special interest group. 2023. Ethics for Incident Response and Security Teams. <https://ethicsfirst.org/>
- [24] Luciano Floridi, Josh Cows, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, et al. 2018. AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds and machines* 28 (2018), 689–707.
- [25] Paul Formosa, Michael Wilson, and Deborah Richards. 2021. A principlist framework for cybersecurity ethics. *Computers and Security* 109 (2021), 102382. <https://doi.org/10.1016/j.cose.2021.102382>
- [26] Dieter Gollmann. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Authentication, Authorisation & Accountability. <https://www.cybok.org/ KA Version 1.0.2>
- [27] Leo A. Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [28] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field methods* 18, 1 (2006), 59–82.
- [29] S Hansson. 2013. *The ethics of risk: Ethical analysis in an uncertain world*. Springer.
- [30] Sven Ove Hansson. 2004. Fallacies of risk. *Journal of Risk Research* 7, 3 (2004), 353–360.
- [31] Jonathan Herington. 2018. The contribution of security to well-being. *J. Ethics & Soc. Phil.* 14 (2018), 179.
- [32] Mireille Hildebrandt. 2013. Balance or trade-off? Online security technologies and fundamental rights. *Philosophy & Technology* 26 (2013), 357–379.
- [33] ISSA. 2023. ISSA Code of Ethics.
- [34] Danish Jamil and Muhammad Numan Ali Khan. 2011. Is ethical hacking ethical? *International Journal of Engineering Science and Technology* 3, 5 (2011), 3–758.
- [35] Tadayoshi Kohno, Yasemin Acar, and Wulf Loh. 2023. Ethical Frameworks and Computer Security Trolley Problems: Foundations for Conversations. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 5145–5162. <https://www.usenix.org/conference/usenixsecurity23/presentation/kohno>
- [36] Wenke Lee. 2019. *The Cyber Security Body of Knowledge v1.0, 2019*. University of Bristol, Chapter Malware & Attack Technology. <https://www.cybok.org/ KA Version 1.0>
- [37] Michele Loi and Markus Christen. 2020. *Ethical frameworks for cybersecurity*. Springer International Publishing.
- [38] Kevin Macnish and Jeroen van der Ham. 2021. Ethical Approaches to Cybersecurity. (2021).
- [39] Mary Manjikian. 2022. *Cybersecurity ethics: an introduction*. Taylor & Francis.
- [40] Andrea M Matwyshyn, Ang Cui, Angelos D Keromytis, and Salvatore J Stolfo. 2010. Ethics in security vulnerability research. *IEEE Security & Privacy* 8, 2 (2010), 67–72.
- [41] Gwenth Morgan and Bert Gordijn. 2020. A care-based stakeholder approach to ethics of cybersecurity in business. *The ethics of cybersecurity* 21 (2020), 119–138.
- [42] Francois Mouton, Mercia M Malan, Kai K Kimppe, and HS Venter. 2015. Necessity for ethics in social engineering research. *Computers & Security* 55 (2015), 114–127.
- [43] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.
- [44] National Institute of Standards and Technology. [n. d.]. Cybersecurity definition.
- [45] Frank Piessens. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Software Security. <https://www.cybok.org/ KA Version 1.0.1>
- [46] Awais Rashid, George Danezis, Howard Chivers, Emil Lupu, Andrew Martin, Makayla Lewis, and Claudia Peersman. 2018. Scoping the cyber security body of knowledge. *IEEE Security & Privacy* 16, 3 (2018), 96–102.
- [47] David G Robinson and J Alex Halderman. 2012. Ethical issues in e-voting security analysis. In *Financial Cryptography and Data Security: FC 2011 Workshops, RLCPs and WECSR 2011, Rodney Bay, St. Lucia, February 28–March 4, 2011, Revised Selected Papers 15*. Springer, 119–130.
- [48] Phillip Rogaway. 2015. The moral character of cryptographic work. *Cryptology ePrint Archive* (2015).
- [49] Christian Rossow and Sanjay Jha. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Network Security. <https://www.cybok.org/ KA Version 2.0.0>
- [50] Vassil Roussev. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Forensics. <https://www.cybok.org/ KA Version 1.0.1>
- [51] M. Angela Sasse and Awais Rashid. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Human Factors. <https://www.cybok.org/ KA Version 1.0.1>
- [52] Clive Seale. 1999. Quality in qualitative research. *Qualitative inquiry* 5, 4 (1999), 465–478.
- [53] Jake M L Stein, Vidminas Vizgirda, Max Van Kleek, Reuben Binns, Jun Zhao, Rui Zhao, Naman Goel, George Chalhoub, Wael S Albayaydh, and Nigel Shadbolt. 2023. You Are You and the App. There's Nobody Else.: Building Worker-Designed Data Institutions within Platform Hegemony. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 281, 26 pages. <https://doi.org/10.1145/3544548.3581114>
- [54] Gianluca Stringhini. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Adversarial Behaviours. <https://www.cybok.org/ KA Version 1.0.1>
- [55] Carmela Troncoso. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Privacy & Online Rights. <https://www.cybok.org/ KA Version 1.0.2>
- [56] United States. 1978. *The Belmont report: ethical principles and guidelines for the protection of human subjects of research*. Vol. 1. Department of Health, Education, and Welfare. National Commission for the Protection of Human Subjects of Biomedical.
- [57] Ibo van de Poel. 2020. Core values and value conflicts in cybersecurity: beyond privacy versus security. *The ethics of cybersecurity* (2020), 45–71.
- [58] Srdjan Capkun. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Physical Layer & Telecommunications Security. <https://www.cybok.org/ KA Version 1.0.1>
- [59] Karsten Weber and Nadine Kleine. 2020. Cybersecurity in health care. *The Ethics of Cybersecurity* (2020), 139–156.
- [60] Laurie Williams. 2021. *The Cyber Security Body of Knowledge v1.1.0, 2021*. University of Bristol, Chapter Secure Software Lifecycle. <https://www.cybok.org/ KA Version 1.0.2>

A ETHICS OF THE CYBER SECURITY PROFESSION: INTERVIEW GUIDE

Scoping Questions

- (1) Can you tell us about yourself? How long have you worked in cyber security?
- (2) What is your role in the company that you work at?
- (3) When did you join the company?
- (4) What are your responsibilities?
- (5) In which of the following areas of cybersecurity have you had any professional involvement:
 - Human, Organisational & Regulatory Aspects (e.g. risk management, compliance, law, human factors, privacy)
 - Attacks & Defences (e.g. malware, forensics, incident management)
 - Systems Security (e.g. cryptography, virtualisation security, Identity/Authentication/Authorization)
 - Software and Platform Security (Software Security, Web & Mobile Security, Software Development Lifecycle)
 - Infrastructure Security (e.g. Network Security, Hardware Security, Cyber Physical Security, Telecommunications Security)

General Ethical Questions

- (1) In your opinion, what are the most important ethical considerations for cyber security professionals?
- (2) Do you think that cyber security professionals should be held to the same ethical standards as other professions?
- (3) What challenges or obstacles do cyber security professionals face when trying to adhere to ethical principles?
- (4) How can we ensure that the cyber security profession is accountable and transparent in its decisions and actions?
- (5) Do you think that cyber security professionals should be held liable for any harm caused by their activities?

Professional Ethical Practices

- (1) How do you ensure that your cyber security activities are conducted in an ethical manner? How do you evaluate the ethical implications of a cyber security decision or action?
- (2) What ethical issues do you encounter in your day-to-day work as a cyber security professional? Have you ever experienced ethical dilemmas in your profession? Can you relate them to the following areas of cybersecurity:
 - Human, Organisational & Regulatory Aspects (e.g. risk management, compliance, law, human factors, privacy)
 - Attacks & Defences (e.g. malware, forensics, incident management)
 - Systems Security (e.g. cryptography, virtualisation security, Identity/Authentication/Authorisation)
 - Software and Platform Security (Software Security, Web & Mobile Security, Software Development Lifecycle)
 - Infrastructure Security (e.g. Network Security, Hardware Security, Cyber Physical Security, Telecommunications Security)
- (3) Have you ever seen or experienced situations where harm occurs to users, employees, or organisations arising from cybersecurity activities or technologies? How would you handle it? (**Non-maleficence**)

- (4) How do you ensure that cyber security activities are conducted with respect for the rights of users and organisations? How do you make sure that your activities are beneficial to humans, promote human well-being, and make our lives better overall? (**Beneficence**)
- (5) How do you ensure that your cyber security activities respect human autonomy? How do you make sure that users use informed decisions for themselves about how that technology is used in their lives? (**Autonomy**)
- (6) How do you ensure that your cyber security activities promote fairness, equality, and impartiality? How do you make sure that they do not unfairly discriminate, undermine solidarity, or prevent equal access? (**Justice**)
- (7) How do you ensure that your cyber security activities are used in ways that are intelligible, transparent, and comprehensible? How do you make sure it is accountable and responsible for its use for cyber security technologies? (**Explicability**)
- (8) What do you think the role and implications of AI will be for cybersecurity ethics?