Zrahia, Aviram

Article

Navigating vulnerability markets and bug bounty programs: A public policy perspective

Internet Policy Review

Provided in Cooperation with:
Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin

Suggested Citation: Zrahia, Aviram (2024) : Navigating vulnerability markets and bug bounty programs: A public policy perspective, Internet Policy Review, ISSN 2197-6775, Alexander von Humboldt Institute for Internet and Society, Berlin, Vol. 13, Iss. 1, pp. 1-30, https://doi.org/10.14763/2024.1.1740

This Version is available at:
https://hdl.handle.net/10419/285315

INTERNET POLICY REVIEW
Journal on internet regulation

Volume 13 **Issue** 1

RESEARCH ARTICLE

OPEN ACCESS

PEER REVIEWED

Navigating vulnerability markets and bug bounty programs: A public policy perspective

Aviram Zrahia *Tel Aviv University*

DOI: https://doi.org/10.14763/2024.1.1740

Published: 15 February 2024
Received: 5 October 2023 **Accepted:** 6 December 2023

Competing Interests: The author has declared that no competing interests exist that have influenced the text.

Citation: Zrahia, A. (2024). Navigating vulnerability markets and bug bounty programs: A public policy perspective. *Internet Policy Review*, *13*(1). https://doi.org/10.14763/2024.1.1740

Keywords: Cybersecurity policy, Bug bounty programs, Economics of vulnerabilities, Digital market, Vulnerability sharing

Abstract: As societies become increasingly dependent on digital means, organisations seek ways to prevent software exploitation by eliminating vulnerabilities or acquiring them as products. However, there is an ongoing debate regarding the extent to which governments should become involved in markets for vulnerability sharing. This paper examines the economics of vulnerabilities and outlines possible areas for governmental interventions. I survey three policy alternatives to support the discovery and disclosure of software vulnerabilities: integrating security and penetration testing into the software development life cycle, acquiring exploitable critical vulnerabilities by governments, and promoting bug bounty programs and platforms as vulnerability-sharing structures. For each suggested alternative, I present an impact matrix to qualitatively measure the effectiveness and efficiency of the vulnerability discovery process and the attractiveness, legality and trustworthiness of the disclosure process. I argue that bug bounty programs that bring together organisations and ethical hackers to trade vulnerabilities produce the highest impact. These gig economy structures are often based on two-sided digital market platforms as their foundation and offer a low entry barrier and assurance level for both market players. The discussion provides a foundation for governmental decision-makers to design effective policies for sharing vulnerabilities.

1. Introduction

Cyber threats emerge as a severe problem for the global economy as society increasingly depends on information technology for all its functions (World Economic Forum, 2023). The actions of offenders can have far-reaching consequences, affecting the financial, social and health well-being of individuals, organisations and nations. Consequently, policymakers have recognised the need to protect citizens and firms against cyber threats, leading to a growing trend in designing, adopting and implementing cybersecurity-related governmental policies. According to Ankit Fadia et al. (2020), more than 100 governments have developed national cybersecurity defence strategies, and some have also established dedicated National Cybersecurity Agencies (NCAs) to help protect the public against cybersecurity attacks, theft, fraud and abuse.[1]

There are various methods that cyber attackers can use; however, exploiting software vulnerabilities (bugs) remains a significant attack vector (*Cyber Attacks Statistics*, n.d.). Common ways to prevent software exploitation before such an attack attempt is made are eliminating vulnerabilities during coding and proactively discovering and fixing them in existing products and services. Leveraging third-party crowd wisdom can make it easier to detect vulnerabilities; I hereafter refer to these vulnerabilities as the product and their trading activity as the market for vulnerabilities.

This article examines the following public policy problem: How should governments become involved in vulnerability-sharing markets? I further focus on the research question: How can policymakers support discovering and disclosing software vulnerabilities in systems, products and services? To address this question, I survey three areas of possible intervention. The first, sometimes called security-by-design, aims to prevent vulnerabilities introduced during development by integrating security into the software development life cycle (SDLC), and it may involve penetration tests executed by internal or contractor teams. The second aims to reduce the number of high-impact exploits in the black market by acquiring highly exploitable critical vulnerabilities directly or through an intermediary entity. The third is promoting bug bounty programs and platforms as mediation entities between individual security researchers and firms. The expected outcome of this alternative is an increased number and quality of software vulnerabilities discovered by the ethical hacker community and disclosed to the public so that they can de-

1. Examples include the National Cyber Security Centre (NCSC) in the UK, and the Cybersecurity and Infrastructure Security Agency (CISA) in the US.

fend against them.

Using a simplified rationalist policy analysis process, I identify two primary goals for the proposed alternatives, which align with the research question. The first goal is to enhance the efficient discovery of vulnerabilities in products and systems, while the second aims to support a legal and trustworthy vulnerability disclosure process. I qualitatively assess three impact categories for each goal to evaluate the suggested policies before presenting the results in a comparative impact matrix. My informed interpretation indicates that bounty programs and platforms are effective and have low barriers to entry for firms and the ethical hacker community.

Therefore, I recommend increased governmental intervention by promoting or requiring these structures based on commercial or community-driven coordinated disclosure initiatives. While I have identified a preferred policy, it is essential to note that the discussed options are not mutually exclusive and can be implemented in parallel, as acknowledged by ENISA (2023). The paper aims to provide a starting point for policymakers yet to engage in this area by listing possible intervention alternatives and justifying additional investments for governments already active in the market for vulnerabilities.

This article proceeds as follows: Section 2 discusses the dynamics of markets for vulnerabilities and bug bounty programs and lists examples of governmental intervention. Section 3 presents the policy design methodology and highlights some of my considerations. Section 4 describes the problem definition, lists related policy design issues and introduces the evaluated governmental policies. Next, Section 5 compares the solution alternatives, their expected outcomes and the associated trade-offs. Finally, Section 6 concludes the discussion and summarises its main implications and limitations.

2. Background: sharing software vulnerabilities

Cyber-related risks are growing and have become one of the most severe global economic risks the world may face over the next decade (World Economic Forum, 2023). Organisational stakeholders can address these risks in various ways, including avoidance, acceptance, mitigation and transfer (Martin-Vegue, 2021). One approach to mitigate the cyber risks associated with software exploitation is to identify and address security vulnerabilities before they are exploited. This approach is aligned with the “identify” and “protect” risk mitigation stages of the Cybersecurity Framework offered by NIST (2018). Discovering vulnerabilities in products and ser-

vices can be assigned to the firm's development or security-testing teams, outsourced to third-party company experts, or delegated to external individual researchers through bug bounty programs in a trend that aligns with the novel idea of crowdsourcing (Akgul et al., 2020).

Reviewing government interventions in cybersecurity and the background and dynamics of the vulnerability-sharing problem domain is required to understand this market development better.

2.1. Government intervention in cybersecurity

In recent years, governments have established agencies dedicated to protecting their assets and citizens against cyber threats. This task often mandates defining new policies or regulations. Still, progress in mitigating cyber risks is challenging, possibly due to conflicting equities, negative externalities, trade-offs between civil liberties, privacy concerns and more (National Research Council, 2014). The policy alternatives presented in this paper illustrate this challenge.

Governmental intervention in cybersecurity is expected to increase in the coming years as cyber protection becomes a regulatory obligation in many sectors. Governments may increase their direct operational involvement in areas considered national priorities, such as Critical Infrastructure Protection (CIP) or Computer Emergency Readiness Team (CERT).[2] Consequently, they can promote regulations that will hold the private sector liable. For example, the US aims to shift the responsibility and liability for cybersecurity away from individuals, small businesses and local governments and onto the organisations providing products and services (The White House, 2023). Similarly, the EU promotes in its NIS2 Directive legal measures on operators of essential services in specific sectors if they fail to take appropriate security measures or follow incident notification rules (Directive 2022/2555, 2022).

Governments and policymakers have long recognised the importance of cybersecurity information sharing as a collaborative effort to enhance cyber-defence (or -adversary) posture by leveraging the broader community's capabilities, knowledge and experience (Zrahia, 2018).[3] The shared information might include threat-centric indicators/objects, best practices and tools and target-related data objects,

2. For example, the US's 2013 National Infrastructure Protection Plan (CISA, 2013) and the United States Computer Emergency Readiness Team (*US-CERT*, n.d.).

3. An individual hacker stopped the global WannaCry ransomware attack in 2017, showcasing the power of the community (MalwareTech, 2017).

namely software vulnerabilities (Libicki, 2015). This paper concentrates on the latter information type.

2.2. The economics of vulnerabilities

A software vulnerability is "a security flaw, glitch, or weakness found in software code" that an attacker could exploit (NIST-CSRC, n.d.). Identifying and fixing software vulnerabilities, commonly known as the vulnerability life cycle, typically begins with the (unintentional) creation of a bug during the coding phase. Unfortunately, an attacker may find and exploit the vulnerability before it is disclosed and a patch developed, resulting in a zero-day (0-day) exploit. However, once the vulnerability is detected and identified by the development team or a security researcher, efforts are prioritised to create and issue a software patch to eliminate the exposure (Bilge & Dumitras, 2012). Extensive research has been conducted on the vulnerability life cycle, including comprehensive overviews by Shahzad et al. (2012) and categorisations of pre- and post-disclosure risk by Rajasooriya et al. (2016). In addition, Ransbotham et al. (2012) summarise the primary pathways to vulnerability disclosure.

Vulnerabilities for sale may be considered a product of the "knowledge economy" created by knowledge-intensive activities and characterised by rapid obsolescence (Powell & Snellman, 2004, p. 199). Furthermore, like other knowledge or data objects, vulnerabilities are non-rival goods that can be used by several parties concurrently. The market for vulnerabilities as products can be described using microeconomics terminology, where organisations generate demand and security professionals supply their expertise and find them. From a supply chain standpoint, finding a vulnerability may be considered a make-or-buy management decision. Williamson (2008) outlines three governance decisions a company may encounter while assessing Transaction Cost Economics (TCE): markets, hybrids and hierarchies. In light of this definition, companies can meet the demand for vulnerabilities with their development and security personnel utilising internal hierarchies. Alternatively, they could use hybrid long-term contracting of specialised companies or embrace a market strategy with skilled individual security researchers with no bilateral stakeholder dependency.

I embrace the view of Ablon & Libicki (2015) and divide the vulnerability market into three categories: legitimate (white), illegal (black) and legal but anonymous (grey).[4] In the white market, buyers and sellers are identified and may legally trade

4. The terms "white", "grey" and "black" are the standard naming of these markets in the cyber world.

vulnerabilities so vendors can fix them. The underground black market is where cybercrime organisations buy exploits, attack services, stolen assets and other illegal products from black-hat hackers. The grey market facilitates the exchange of vulnerabilities and exploits that might be used for offensive purposes. While this market is not illegal *per se*, it operates in a moral and ethical grey area due to the potential for harm associated with undisclosed vulnerabilities.

The analysis requires an understanding of the way the value of a vulnerability changes depending on its life-cycle stage and traded market (Figure 1). A zero-day vulnerability may be valued at six or more figures in the white and black markets (*Apple Security Bounty Categories*, n.d.; Perlroth, 2021), but its price declines differently over time. In the white market, disclosed vulnerabilities are shared as public goods for free, so their value drops to zero once the vendor releases a patch and they become public goods. In contrast, in the black market, the exploit code has a monetary value even after N-days due to product exclusivity. Regardless, its value drops over time as the likelihood of finding and exploiting a non-patched system decreases.

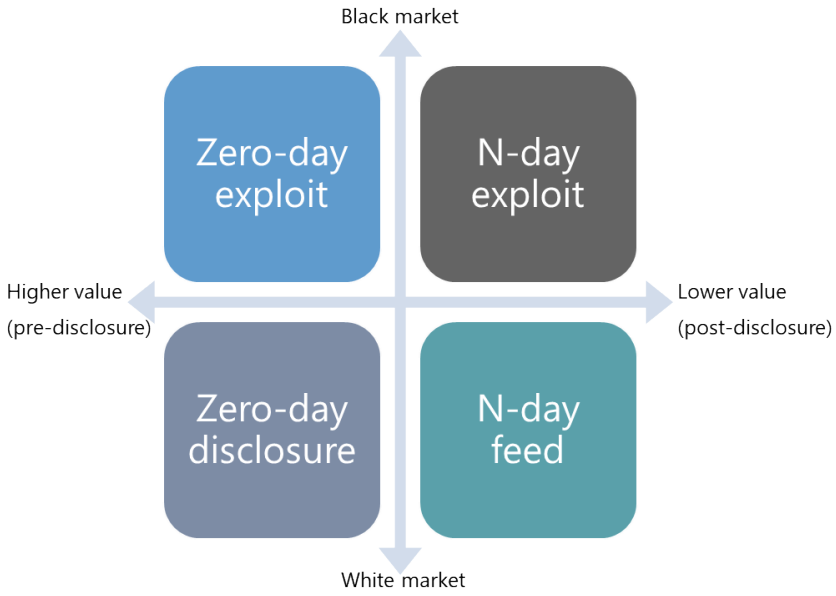

FIGURE 1: Vulnerability as a product value matrix in black and white markets.

2.3. Bug bounty programs and platforms

The claim dubbed 'Linus law', that "given enough eyeballs, all bugs are shallow" (Raymond, 1999, p. 29) refers to the co-development and testing of open-source software involving many people who deliver a less buggy (and therefore less vulnerable) code together. A similar principle may apply to bug bounty programs util-

ising crowdsourced vulnerability discovery. From an organisational perspective, effectively identifying vulnerabilities is a high-value challenge, so firms can benefit from engaging large crowds of researchers to tackle this task. This choice aligns with theories of firm boundaries, which involve deciding which assets, activities and resources to "own" and which to access through the market (Zenger et al., 2011, p. 95).

Bug bounty programs are structured arrangements between organisations and individual security researchers to trade vulnerabilities as products. They allow organisations to interact with cyber-security experts whose knowledge complements the capabilities of the firm's development and testing teams. Through this exchange, security researchers can report on security vulnerabilities and receive legitimate compensation for their findings and recognition from their peers and the industry for their expertise (Bienz & Juranek, 2020; Malladi & Subramanian, 2020).

Bug bounty platforms are two-sided digital marketplaces that host multiple bug bounty programs, bringing together security researchers and organisations to facilitate vulnerability trading (Maillart et al., 2017; Subramanian & Malladi, 2020; Wachs, 2022; Zhao et al., 2017). These platforms reward the first participant submitting a novel vulnerability report with a direct or indirect payment (bounty), creating a tournament-like arrangement (Jo, 2020). Using these platforms reduces information asymmetries and other frictional costs associated with the transaction of specific, infrequent, and uncertain assets (Wachs, 2022). Figure 2 illustrates the role of a bug bounty platform as a facilitator for the vulnerability-sharing transaction between organisations as buyers and researchers as sellers.

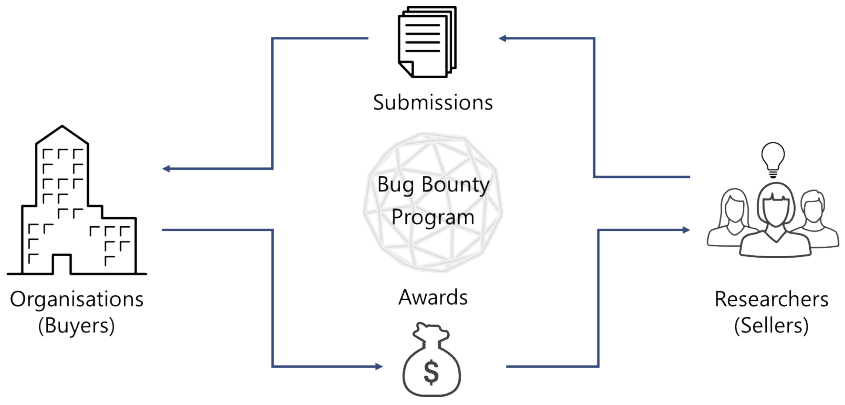

FIGURE 2: A bug bounty platform as a two-sided market for sharing software vulnerabilities.

Bug bounty programs can be viewed as a competitive economy model where buyers and sellers attempt to maximise their utility and profits. The individual re-

searchers are the sellers, the organisations generating the bounty programs are the buyers and the discovered vulnerabilities are the goods.[5] These structures are also related to the phenomenon of the gig economy: a labour market for independent contracting that happens through, via and on digital platforms (James, 2021). In this setting, a vulnerability report constitutes a mere gig economy transaction with a low entry barrier, potentially allowing policymakers to control the supply size (Zrahia et al., 2022).

3. Policy design methodology

Policy analysis involves explaining problems related to the general public and developing alternatives to address them and mitigate their failures. I followed the simplified rationalist process of Weimer & Vining (2017) by dividing the effort into two main components. Section 4 focused on problem analysis, which includes defining the problem, listing potential solutions and setting policy goals. In Section 5, I conducted a solution analysis, further expanding on the policy options while predicting, evaluating and comparing their impacts. I began the problem analysis with the research question: How can policymakers support discovering and disclosing software vulnerabilities in systems, products and services? Next, I argued that the problem is a societal concern that justifies governmental intervention, and I listed the three evaluated solutions. Finally, I briefly described how each alternative would help and its expected outcome.

I selected two goals, viewing the first as "substantive" and the second as "procedural" (Bali et al., 2021). The impacts associated with the substantive goal are directly concerned with the ends of the policy, while the results related to the procedural goal, indirectly but significantly, affect processes and outcomes accounting for the means to achieve the policy ends. Furthermore, I set the goals so that each represents a different stage in the vulnerability-sharing process (discovery and disclosure) and a different stakeholder perspective in the market for vulnerabilities — the first goal was associated with the buyer (organisation) perspective, while the second pertained to the seller (researcher).

I detailed the three policy alternatives in the solution analysis stage, referencing relevant academic literature and professional resources. To measure the impact of each policy on each goal, I used qualitative scoring criteria of three levels: low, medium and high. These categories are not necessarily equally spaced, and I used ranges within these classes where appropriate. I presented the results in a com-

5. This viewpoint refers to the market after a researcher finds a vulnerability (ex-post).

parative impact matrix that includes the essence of the analysis. Then, I detailed the trade-offs, tensions and misalignments between the market players associated with each choice. Finally, the policy advice/recommendation reflects the impact matrix after qualitatively comparing the results.

The analysis is subjective and influenced by my informed interpretation.[6] Hence, I do not claim to provide a precise objective measure of policy superiority. Furthermore, while there is a growing interest in methods for qualitative comparative analysis in public policy and other disciplines (Brans & Pattyn, 2017; Fischer & Maggetti, 2017), I have chosen a different approach for two reasons. First, scoring and assigning (sometimes arbitrary) weights to the impact categories would still be subjective and, therefore, keep the essence of the analysis the same. The second reason for avoiding a complex scoring methodology is that the interventions evaluated are complementary and could be promoted in parallel. For example, bug bounties can provide continuous software testing against vulnerabilities and become an integrated part of the SDLC process (Bugcrowd, 2018). Instead, I followed Weimer & Vining (2017) recommendation to use detailed comparisons of the alternatives rather than decision rules, as the latter tend to divert attention from the trade-offs and obscure the outcomes of the suggested options.

In addition, some may question the practicality of using a rational-based policy approach to address issues in the diverse cybersecurity problem domain. However, there are two counterarguments to this claim. Firstly, the selected policy alternatives can be implemented in small incremental changes while "muddling through" (Lindblom, 1959, 1979). Secondly, the technical practices that support these policies are already in use to some extent by various organisations and governments, making it easier to assess and compare their impact.

I followed several strategies to enrich the trustworthiness of this qualitative research (Yadav, 2022). First, I synthesised many academic and professional sources, theories and viewpoints. While this may be perceived as a confusing approach, it strengthens the credibility of the discussion by triangulation. I also provided plentiful descriptions to highlight the study's context and ease the transferability of the debate to the reader's surroundings. Finally, I offered an extensive audit trail and references for the arguments presented to increase the paper's dependability.

6. As a cybersecurity practitioner and an active member of ISACA, ISC2, and SANS professional organisations, I am a true believer in cyber information sharing and the power of individuals and communities to help society better address cyber challenges.

4. Problem analysis

Hereafter, I refine the policy-related question and survey three governmental intervention alternatives for sharing vulnerabilities, their goals and respective impact outcomes.

4.1. Problem definition

I define the following public policy problem: How should governments become involved in vulnerability-sharing markets? I further focus on the research question: How can policymakers support discovering and disclosing software vulnerabilities in systems, products and services? The vulnerability-sharing process is a subset of cybersecurity information sharing (Libicki, 2015; Mermoud et al., 2019), and the research question covers the two stages that make it valuable for mitigating software vulnerability risks (Rajasooriya et al., 2016).

The first stage, vulnerability discovery, can be achieved through internal or outsourced secure development and penetration testing practices (Jones & Rastogi, 2004) or by utilising bug bounty programs and platforms (Malladi & Subramanian, 2020). I consider both approaches as two of the analysed alternatives. The second stage, disclosure, refers to the process of revealing the existence of the vulnerability, first to the vendor and eventually to the public. Empirical analysis suggests that disclosure can delay and lower the likelihood of an initial attack, limit its spread and decrease its total instances (Ransbotham et al., 2012). However, not all vulnerabilities are disclosed or shared with the public through black or grey market transactions (Ablon & Libicki, 2015; McKinney, 2007).

4.2. Should governments get involved?

Governments can arguably "do more to encourage and accelerate the process of finding software vulnerabilities with modest amounts of funding and without passing new legislation" (Libicki, 2015, p. 3). However, before considering government involvement, it is essential to determine whether the issue at hand is a societal concern requiring public resources.

Market failures "provide the traditional economic rationales for public participation in private affairs" (Weimer & Vining, 2017, p. 74). They might occur if a product or service has positive externalities affecting the general public and the allocation of resources is not optimal. For example, a situation where it is difficult to collect payment from all potential beneficiaries of a good or service (Bardach & Patashnik, 2020) might lead to such a failure, called the "free rider" problem

(Hardin & Cullity, 2020; Varian, 2004). Indeed, in the market for vulnerabilities, seller or buyer actions can impact third-party entities that do not actively participate in the two-sided program transaction. Böhme (2006) mentions two possible reasons for market failures in the computer security market. One is the "lemon market problem" (Akerlof, 1970) — where consumers cannot distinguish between secure and insecure products, leading to lower prices — and the other is the "free rider" dilemma discussed above.

Furthermore, software quality is a public concern (IEEE-CS/ACM Joint Task Force on Software Engineering Ethics and Professional Practices, n.d.), even if it is owned by private companies and consumed directly by only part of the public. Modern products and services have multiple integrated components, and attacks can propagate between systems and organisations in this digital ecosystem. Hence, a compromised asset in one place can lead to a supply chain attack that affects many organisations (Ohm et al., 2020). A visible example is the 2020 SolarWinds incident (Peisert et al., 2021), which impacted numerous governmental and commercial organisations worldwide.

I follow the flowchart Weimer & Vining (2017) suggest for linking market and government failures to policy interventions. First, I consider whether an "operational market" for vulnerabilities exists, reflecting a significant, legally priced and accessible formation. While a white market for vulnerabilities does exist, the grey and black markets are more lucrative (Libicki et al., 2015). Therefore, government intervention might be in place to further operationalise the market and diminish illegal transactions with possible negative externalities.

4.3. Evaluated alternatives

Table 1 summarises the three evaluated policies for governmental intervention. The first policy encourages secure development and penetration testing and proactively prevents bugs during product development. The second policy involves acquiring zero-day vulnerabilities and disclosing them to the vendor and the public, resulting in fewer high-impact exploits in the black market. The third policy promotes bug bounty programs and platforms, which should increase the safe disclosure of vulnerabilities by the community of ethical hackers.

The policies examined in this study are by no means the only possible governmental interventions in the cybersecurity space. They were selected to meet three criteria. Firstly, they all aim to address vulnerabilities discovery and disclosure prior to the occurrence of a cybersecurity incident. Measures taken to mitigate known

vulnerability risks, such as preventing exploitation of vulnerabilities or patching vulnerable systems, are outside the scope of this analysis. Secondly, organisations and governments can implement these policies using small incremental steps, as Lindblom (1959) suggested. This is important as the simplified rationalist approach may be criticised for being unrealistic in the complex problem domain of cybersecurity. Finally, the technical practices underlying these policies are already in use to some extent by organisations and governments, making it easier to assess and compare their impact.

While the primary expected outcome of each policy is highlighted, it may also apply to other evaluated alternatives. In addition, using these policies concurrently is consistent with the best practices outlined in the literature (Walshe & Simpson, 2022) and the cybersecurity approach of a layered defence.[7]

TABLE 1: Methods and outcomes of governmental intervention policy options related to sharing vulnerabilities

GOVERNMENTAL INTERVENTION	METHOD	EXPECTED OUTCOME
Encourage secure development and penetration tests	Influence vulnerability self-discovery by regulation or recommendation	Prevent software vulnerabilities during development
Acquire zero-day vulnerabilities	Utilise existing grey-market intermediaries	Reduce the number of high-impact exploits available in the black market
Promote bug bounty programs and platforms	Encourage or enforce the use of bug bounty programs and platforms	Increase the number and quality of vulnerabilities safely discovered and disclosed by the community

4.4. Goals and impacts

The proposed alternatives have two primary goals and three impact categories for assessment, as outlined in Table 2. Each goal targets a different market participant; balancing them is imperative to operationalise the market further despite the inherent tension.

TABLE 2: Policy goals and impacts for governmental intervention related to sharing vulnerabilities

GOALS	Effective and efficient discovery of vulnerabilities	An attractive, legal, and trustworthy disclosure process
IMPACTS	1. Number of unique discovered vulnerabilities	1. Ease of reporting
	2. Quality of discovered vulnerabilities	2. Safe harbour disclosure
	3. Economic efficiency in production	3. Trustworthy and attractive model

7. Using multiple (layered) security solutions is a strategy deployed with the goal that a subsequent defence layer will stop an attacker if they succeed in penetrating one layer of defence.

My primary objective is to assist organisations in effectively and efficiently discovering vulnerabilities. To evaluate this substantive goal, I employ three impact measurements directly concerned with the ends of the selected policy (Bali et al., 2021). The first is the number of unique and valid vulnerabilities discovered without submissions incorrect or already known to the organisation. The second measurement is the quality of discovered vulnerabilities, which reflects their value to the buyer. Lastly, I consider economic efficiency, accounting for purchasing and processing costs.

My second goal is to facilitate an attractive, legal and trustworthy vulnerability disclosure process. While this goal primarily aligns with the researcher's interests, it should encourage all market players to share vulnerabilities. I use three impact measurements associated with this procedural goal to indirectly impact processes and outcomes to achieve the first goal's policy ends (Bali et al., 2021). The first is the ease of reporting, reflecting the usability of the submission process and any barriers to entry. The second is compliance with the set of legal requirements known as "safe harbour" that allows researchers acting in good faith "to provide security feedback without fear of legal repercussions" (*The Disclose.Io Project*, n.d.). Lastly, I consider whether the policy incentivises researchers and supports a trustworthy and attractive model acceptable to all stakeholders.

5. Solution analysis

5.1. Solution details

Below, I explain each of the suggested solutions and discuss their benefits along with relevant caveats.

5.1.1. Encourage secure development and penetration tests

A secure software development life cycle (SDLC) framework allows organisations to promote continuous and iterative software security through a structured approach spanning planning, development, testing, deployment and maintenance. Embedding security into SDLC requires incorporating it into all life-cycle stages (McGraw, 2004). This alternative solution combines two best practices: integrating security measures into the SDLC process of designing and building high-quality software and conducting regular penetration tests and code audits to identify and address vulnerabilities.[8] Ramirez et al. (2020) survey and compare secure SDLC

8. Penetration tests involve proactive attempts of internal or external security researchers to exploit vulnerabilities and hack the system, while a code audit examines the source code looking for vulnerabilities.

standards, guidelines and certifications. A best practice example supporting this alternative is the Secure Software Development Framework (SSDF), published by the National Institute of Standards and Technology (NIST), which provides recommendations for mitigating the risk of software vulnerabilities (Souppaya et al., 2022).

Some policymaker initiatives already encourage or enforce these practices. One of the requirements of the Payment Card Industry Data Security Standard (PCI-DSS) for credit card processing is to "develop and maintain secure systems and applications" (PCI Security Standards Council., 2010). Another example is the EU's Cyber Resilience Act (2022), which regulates cybersecurity requirements for products with digital elements to ensure that manufacturers improve their security from the design and development phase and throughout the whole life cycle. As liability for insecure products and services increasingly falls on their owners, governments may continue to promote and enforce these practices.

5.1.2. Acquire zero-day vulnerabilities

Under this policy option, a government may buy, directly or through a grey market intermediary, a highly exploitable zero-day vulnerability or purchase an adversary tool that uses it.[9] The grey market is characterised by anonymity, high monetary gains and the potential to acquire offensive attack tools (Annu-Essuman, 2014). Governments buying vulnerabilities in this market may eventually share them with the affected vendor or keep them private (Ablon & Bogart, 2017).

There is an apparent trade-off between disclosing vulnerabilities to the vendor and thus allowing them to be fixed or retaining them for national security purposes (Schwartz & Knake, 2016). A critical zero-day vulnerability can be used by governments as a cyber weapon against enemies of the state or even internally against public figures in certain circumstances.[10] However, keeping this vulnerability unknown to others also means that the state's own governmental and private systems will not be patched and protected against it, so there is a risk that cybercriminals or other countries' state actors will discover and use it. The US federal government has acknowledged this dilemma and released a policy that determines on a case-by-case basis how the government should treat zero-day vulnerabilities (The White House, 2017).

9. Zerodium is an example of a zero-day broker, and the NSO Group is an example of a developer of a surveillance technology that might include zero-day vulnerabilities.

10. For example, several EU governments used the Pegasus spyware software against journalists, politicians, officials and other public figures, leading to an investigation by the European Parliament (European Parliamentary Research Service, 2023).

Governments not already participating in zero-day markets may consider setting goals and establishing evaluation processes, while those already involved may consider increasing their involvement. Reducing the number of available zero days would reduce the number of cyber-criminals and the state programs that depend on them (Maurer, 2017). However, this depends on whether policymakers want to "drain the swamp" of vulnerabilities or use them for offence.

5.1.3. Support bug bounty programs and platforms

In recent years, internet governance and digital platform regulation have become hot topics for scholars and practitioners (Epstein et al., 2016; Flew & Martin, 2022; Fuster Morell, 2022). This involvement expands beyond social networks and commodity markets into two-sided markets for vulnerabilities. In this alternative, governments may encourage or enforce vulnerability disclosure programs (VDPs) or bug bounty programs (BBPs) in specific vertical segments. The first program type allows researchers to safely submit their reports to organisations without receiving cash rewards, and the latter offers monetary awards for unique (unknown) valid discoveries (Walshe & Simpson, 2022).

Various national-level initiatives have been implemented to facilitate coordinated vulnerability disclosure (CVD) policies. Examples include the US requirement from federal agencies (BOD 20-01: Develop and Publish a Vulnerability Disclosure Policy, 2020), the EU's CVD policy (ENISA, 2022) and the UK's vulnerability disclosure toolkit (The National Cyber Security Centre, 2020). These initiatives are often based on commercial bug bounty platforms that outline discovery and disclosure procedures as part of their program's scope and code of conduct.[11]

Organisations operating bug bounty programs often fail to convey all the formal constraints applicable to hackers, requiring them to understand the laws underpinning safe and legal security research (Walshe & Simpson, 2023). Crowdsourcing security as a service through bug bounty platforms can enable this process safely and legally. Choi et al. (2010) found bug bounty programs to be a welfare-improving policy instrument since they either do not affect the firm's disclosure policy or facilitate a change from non-disclosure to disclosure. Alternatively, governments can support community-driven vulnerability disclosure projects such as Disclose.io, which aims to make vulnerability disclosure safe, simple and standardised for everyone.[12]

11. For example, CISA (the US initiative) uses BugCrowd and EnDyna as their bug bounty platform provider (Goldstein, 2021), while the UK uses HackerOne (Ministry of Defence, 2020).

12. The project provides a comprehensive list of known bug bounty and vulnerability disclosure pro-

In addition, bug bounties promote public transparency by facilitating the disclosure process to the public. When ethical security researchers discover vulnerabilities, they can reveal them through full or coordinated vulnerability disclosure methods (Maillart et al., 2017). Full disclosure pressures software owners to fix the issue immediately, as it involves alerting the public directly. Coordinated vulnerability disclosure, on the other hand, allows vendors to address the vulnerability before sharing the details publicly. The optimal disclosure approach remains a topic of debate (Arora & Rahul, 2005; Choi et al., 2010). However, not all vulnerabilities are disclosed or shared with the public during black or grey market transactions (Ablon & Libicki, 2015; McKinney, 2007).

A viable governmental intervention policy may legally enforce bug bounty programs on specific industry verticals (Zhao et al., 2017), establish joint initiatives with existing bug bounty platforms or support vulnerability-sharing community efforts.

5.2. Policy impacts analysis

In the upcoming sections, the policy goals and their associated impact categories will be qualitatively assessed for each solution alternative, with justifications for the analysis.

5.2.1.Evaluating effective and efficient discovery of vulnerabilities

The substantive goal of an effective and efficient vulnerability discovery is aligned with the organisational interests, and the success of a policy in meeting it can be measured by the number and quality of unique vulnerabilities discovered, as well as its economic efficiency.

I reflect on insights from the open-source literature and argue that the number of unique discovered vulnerabilities depends on the number of researchers, their expertise and their access level to the product or service (Schryen & Kadura, 2009). In-depth vulnerability research is made possible with code-level access rather than treating the system as a black box (McGraw, 2004). Although secure SDLC policies allow for thorough penetration testing, I argue that the number of discovered vulnerabilities may depend on a limited number of security experts compared to the potentially large crowd of individual researchers accessible through bug bounty programs and platforms (Maillart et al., 2017). Hence, I anticipate a moderate number of vulnerabilities to be discovered. Acquiring zero-day vulnerabilities is a

grams, detailing where to submit reports and their respective “safe harbour” status.

strategy typically reserved for selected, highly exploitable cases, a subset of the limited supply of these vulnerabilities (Maurer, 2017). Bug bounty programs can lead to the discovery of a moderate to high number of vulnerabilities (Walshe & Simpson, 2020; Zrahia et al., 2022), depending on the program's rules, incentive structure and degree of openness. The latter refers to the choice between a program that is available to everyone (public) or only to a group of researchers (private) who may be pre-selected and possibly granted elevated access rights to research the product or service (Wachs, 2022).

The second impact category I evaluate is the quality of discovered vulnerabilities, which, similarly to the quantity, is affected by the researchers' expertise and access level. Therefore, I argue that software development and penetration test teams, granted elevated access rights, can find medium-high severity vulnerabilities. Zero-day vulnerabilities traded on the grey market are often of exceptionally high quality due to their exploitable nature (Meakins, 2019). In contrast, the quality of submissions to bug bounty programs can vary between low and high (Walshe & Simpson, 2020; Zrahia et al., 2022), depending on the program's characteristics and tournament structure, facilitating competition among researchers.

Finally, assessing the economic efficiency aspect of a vulnerability discovery policy requires a consideration of its costs and benefits. Secure SDLC and penetration tests are typically paid for by contract according to an agreed Scope of Work (SOW) rather than performance-based payment (Engin, 2023). Hence, I mark its effectiveness level as a medium. By contrast, the cost of purchasing a single zero-day to governments may be extremely high, depending on the vulnerability's severity, the exploit's complexity and how long the vulnerability remains undisclosed (Ablon & Libicki, 2015). The latter factor reflects whether the product (the zero-day vulnerability) is a private good, defined by rivalry in consumption and excludability in ownership and use (Weimer & Vining, 2017). Disclosing the vulnerability to the public or other buyers may affect the cost-effectiveness of this policy option if the designated use of the purchased exploit is offensive. Therefore, the impact category for buying zero-day vulnerabilities may vary between medium and high, assuming the governmental goal is defensive. The economic efficiency of bug bounty programs can be measured using the unique-to-total submission ratio,[13] which considers the effort and cost of processing duplicate or incorrect vulnerability reports (Zrahia et al., 2022).[14] Though invalid reports may significantly burden par-

13. The unique-to-total submission ratio represents the percentage of unique vulnerabilities found out of the total number of submitted reports.

14. Duplicate submission is a discovery of a vulnerability already known or identified by another re-

ticipating organisations (Zhao et al., 2017), there are ways to reduce them by limiting access to the program, changing the rewards structure and more. Therefore, the economic efficiency of bug bounty programs that embrace performance-based payments might be considered medium-high.

5.2.2.Evaluating the disclosure process

The procedural goal of establishing an attractive, legal and trustworthy disclosure process pertains to researchers discovering vulnerabilities more than the organisations acquiring them. The first impact category I suggest for measuring this goal is the ease of reporting which also reflects its barriers to entry. Under the secure SDLC and penetration testing policy the internal workforce or contractors can report vulnerabilities promptly and effectively to the organisation. However, as acknowledged by Çetin et al. (2018), implementing this option requires resource investments and expertise which may challenge small organisations, resulting in a medium barrier to entry. Reporting and barriers to entry are notably more difficult in the grey market for zero-day vulnerabilities, as researchers may lack the necessary connections to sell directly to governments. Hence, introducing an intermediary may facilitate a reporting procedure while preserving anonymity for both parties. Bug bounty programs and platforms in the regulated white market have a more straightforward reporting process based on their predefined scope and rules of engagement.[15] Furthermore, these platforms have relatively low entry barriers since they support a simple registration process for both market players and may allow anyone to submit vulnerability reports.

The second impact measures whether the disclosure process protects the researcher from legal consequences. The SOW for outsourced penetration testing should include clauses that provide legal protection for security researchers. Similarly, the SDLC development process inherently protects employees when fixing bugs. In the grey market policy, researchers selling a zero-day to a government may prefer to remain anonymous and use an intermediary to avoid revealing their identity to the buyer. Moreover, the buying government often wants to maintain the same level of anonymity (Annu-Essuman, 2014). By contrast, bug bounty programs and platforms should include clear disclosure guidance and often support full or partial safe harbour policies to encourage reporting.[16]

searcher. Therefore, it has no value to the vendor (or even negative value considering the costs associated with processing it).

15. For example, Bugcrowd's reporting process (Bugcrowd, n.d.).

16. For example, Microsoft's safe harbour policy (Microsoft, n.d.).

The last impact analysed is a trustworthy and attractive model. Trust between buyers and sellers is associated with better exchange performance, lower transaction costs and enhanced knowledge transfer (Poppo et al., 2016). In the vulnerability market, sellers face an additional challenge related to trust: they must prove the authenticity of the vulnerability without revealing it to the buyer.

Internal or external organisational teams involved in secure SDLC and penetration tests get paid by contract, making trust a non-issue and the attractiveness of the transaction negotiable. However, trust is a real challenge to both sides when selling zero-days to governments directly or through a grey market broker. Researchers may require anonymised cryptocurrency payments and assurance that they will be fulfilled. On the other hand, buyers who want to obtain a zero-day for offensive purposes need exclusive access and non-disclosure commitments to maintain its value. Therefore, using intermediaries may add trust validation and verification to the transaction as Ablon & Libicki (2015) noted. Similarly, trust is critical in the white vulnerability-sharing market where bug bounty programs and platforms facilitate interactions between two entities that may not have any pre-existing relationship or history of interaction. Platform intermediaries can reduce the risk for both parties by ensuring mutually beneficial terms and conditions for disclosure and participation (Subramanian & Malladi, 2020). A coordinated vulnerability disclosure process grants the vendor the necessary time to apply a patch before sharing the vulnerability with the public. No exclusivity risk exists if the organisation intends to fix the issue and notify the public. Vendors usually complete their payments to researchers to maintain their reputation. However, a two-sided bug bounty platform can reduce the risk of unrelieved contractual hazards and add trust and assurance to timely payments.

5.2.3. Comparative impact matrix

The comparative impact matrix presented in Table 3 summarises the predicted impact of each alternative on the two defined goals.

TABLE 3: Impact evaluation of policy options for governmental intervention related to sharing vulnerabilities

IMPACT CATEGORY	POLICY ALTERNATIVES		
	Encourage secure development and penetration tests	Acquire zero-day vulnerabilities	Support bug bounty programs and platforms
GOAL I: EFFECTIVE AND EFFICIENT DISCOVERY OF VULNERABILITIES			
NUMBER OF UNIQUE DISCOVERED VULNERABILITIES	Medium, given the trade-off between thorough testing and the number of security researchers	Low, as only zero-day vulnerabilities are potentially purchased	Medium-high, depending on the program's characteristics
QUALITY OF DISCOVERED VULNERABILITIES	Medium-high, based on the defined scope, access level and expertise of the testing team	High, as it pertains to zero-day critical vulnerabilities only	Varies between low and high, based on the program's characteristics
ECONOMIC EFFICIENCY IN PRODUCTION	Medium, as the penetration task paid regardless of the findings	Varies between medium to high and affected by exclusivity risk and extreme prices of premium vulnerabilities	Medium-high, taking into consideration the valid-to-total submission ratio
GOAL II: AN ATTRACTIVE, LEGAL AND TRUSTWORTHY DISCLOSURE PROCESS			
EASE OF REPORTING AND BARRIERS TO ENTRY	High (easy reporting), as defined in the scope of work of the internal workforce or contractors	Low (complicated reporting), unless made through a third party, grey market broker	High (easy reporting), based on the intermediary platform tools and procedures
LEGALLY SAFE DISCLOSURE	High (legally safe) as defined explicitly in the scope of work for penetration tests and inherent to the SDLC process	Medium, due to the risks associated with grey market transactions	High (legally safe), based on the bounty program's full or partial safe harbour policy
TRUSTWORTHY AND ATTRACTIVE MODEL	Highly reliable payment per contact	Low-medium reliability for both sides, as payments might be conditional and exclusivity questionable	Payment is highly reliable as defined in the program's scope and platform rules

5.2.4. Trade-offs, tensions and misalignments

This section highlights potential areas of misalignment among different market

players or policymakers related to the goals and impact categories. These conflicts may generate market failures, suggesting that public policy oversight is essential.

Firstly, I consider the trade-off vendors face between the quantity and quality of discovered vulnerabilities. It is widely accepted that finding high-severity vulnerabilities requires more expertise and is therefore less common than finding low-quality ones. Additionally, the access level of researchers to the code can impact their ability to find bugs (Schryen & Kadura, 2009). While more researchers testing the code increase the number of discovered vulnerabilities (Maillart et al., 2017), high volumes of low-quality reports can burden operators and consume resources (Walshe & Simpson, 2022). Furthermore, contests can encourage innovation, but admitting more competitors can create tension between innovation and incentives for all players (Terwiesch & Xu, 2008). While having more competitors may stimulate innovation by a higher likelihood that at least one agent will find a highly valued solution, it can also reduce the expected reward of researchers and their incentive to report vulnerabilities. Additionally, attracting high-quality researchers and finding high-severity bugs is becoming more difficult as the program matures, hence the recommendation to increase rewards over time (McCracken, 2019). In light of these issues, I argue that organisations could evaluate their strategy over time and rebalance the number of security researchers and their access level to the code with bug bounty platforms potentially helping match researchers with relevant expertise to specific programs (Kestelyn & Bugcrowd Head of Product Marketing, 2022).

Next, I consider the offensive versus defensive use of zero-day vulnerabilities by governments operating in the grey market. The value of a zero-day primarily depends on its scarcity and secrecy (Meakins, 2019). While buying vulnerabilities may be cost-effective if shared with the public some governments may prefer to keep them private for offensive use. Under the latter scenario, they risk losing exclusivity to the vulnerability. If that happens, its value will decline or completely diminish once patched. The governmental decision to keep or reveal a vulnerability may also depend on the geopolitical circumstances. Furthermore, a concern associated with this alternative is whether the government can be trusted to implement it for the benefit of the general public. The Electronic Privacy Information Center (EPIC) criticises the Vulnerabilities Equities Policy of the US mentioned earlier for lack of transparency, privacy implications, and more (The Electronic Privacy Information Center, n.d.). Security researchers may need to balance the expected reward from a vulnerability discovery and the consequences of illegal, immoral or unethical submission to the black or grey markets. These markets might pay ten times (or more)

higher than the white market would pay, so a researcher who finds a zero-day faces a significant dilemma (Ablon & Libicki, 2015). The procedural goal of facilitating an attractive, legal and trustworthy disclosure process aims to balance this tension and solicit white market transactions.

Another aspect to consider is the applicability of the alternatives in a diverse socio-economic context. The evaluated policies require cooperation with private organisations and the general public and may be affected by different scenarios. During the COVID-19 pandemic for example, there was a huge increase in individual researchers' participation and vulnerability submissions in bug bounty platforms (Zrahia et al., 2022), reflecting the gig-economy nature of this policy option and its sensitivity to external shocks.

Finally, the trust challenge arises when there is no bilateral dependency between the seller and the buyer. Intermediating platforms can solve this tension, as illustrated by various e-commerce, knowledge economy and sharing economy literature sources (Akhmedova et al., 2021; Soleimani, 2022; Zanini & Musante, 2013). In the bug bounty programs setting, platforms can act as a trusted third party to ensure that the interests of both parties are met (Miller, 2007). They can help researchers prove the validity of discovered vulnerabilities without disclosing them and help organisations verify the validity of vulnerabilities before making payments.

6. Concluding thoughts and limitations

Governments and other policymakers have become highly concerned with protecting the cyber domain and their involvement in this space is growing. Governmental intervention in the market for vulnerabilities may shift transactions from the black market to the white market and improve the security posture of systems, products and services used by the public directly or indirectly. Some governments have already implemented related policies, but others are less involved in this market. Well-considered and carefully thought-out policies can provide valuable oversight of sharing initiatives, advance vulnerability identification and maximise social welfare.

This paper explores three possible alternatives for governmental intervention in the market for vulnerabilities: implementing secure development and penetration tests, acquiring zero-day vulnerabilities and supporting bug bounty programs and platforms. I present an impact matrix qualitatively measuring the goals associated with the discovery and disclosure processes for each potential intervention area.

While the alternatives are not mutually exclusive, bug bounty programs and platforms produce the highest impact and have a relatively low barrier to entry for both the organisations and the ethical hacker community.

Therefore, I advise governmental entities, regulation authorities and other policy decision-makers to consider, encourage or prescribe bug bounty programs and platforms based on commercial or community-driven coordinated disclosure initiatives. Despite the extensive literature on the vulnerability market, this paper compares relevant governmental intervention alternatives from a public policy perspective, highlighting their trade-offs, tensions and misalignments, thus contributing to the problem domain. Discussing these structures may serve as a starting point or guideline to motivate the design of a detailed governmental policy. However, further investigation of researcher incentives, firm motivations and bug bounty platform strategies is required to design more effective programs. A follow-up study may check how bug bounty platform operators could attract more buyers and sellers to encourage vulnerability sharing, and organisations maximise their utility function from bug bounty programs.

The limitations of the paper are twofold. First, the recommendation is subjective and reflects my informed interpretation, as there is no precise objective measure of policy superiority. Second, though the paper compares three policies for governmental intervention, evaluating and comparing instruments needed to implement the selected option is beyond its scope and requires more exploration.

ACKNOWLEDGEMENTS

I am grateful to Alon Tal for insightful discussions and continuous guidance. I also thank the reviewers for their constructive feedback.

References

Ablon, L., & Bogart, A. (2017). *Zero days, thousands of nights: The life and times of zero-day vulnerabilities and their exploits*. RAND Corporation. https://doi.org/10.7249/RR1751

Ablon, L., & Libicki, M. (2015). Hackers' bazaar: The markets for cybercrime tools and stolen data. *Defense Counsel Journal*, *82*(2), 143–152. https://doi.org/10.12690/0161-8202-82.2.143

Akerlof, G. A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, *84*(3), 488. https://doi.org/10.2307/1879431

Akgul, O., Eghtesad, T., Elazari, A., Gnawali, O., Grossklags, J., Votipka, D., & Laszka, A. (2020). The hackers' viewpoint: Exploring challenges and benefits of bug-bounty programs. *Proceedings of the 6th Workshop on Security Information Workers (WSIW 2020)*. Workshop on Security Information Workers (WSIW 2020). https://www2.cs.uh.edu/~gnawali/papers/bugbounty-wsiw20-abstract.html

Akhmedova, A., Vila-Brunet, N., & Mas-Machuca, M. (2021). Building trust in sharing economy platforms: Trust antecedents and their configurations. *Internet Research*, *31*(4), 1463–1490. https://doi.org/10.1108/INTR-04-2020-0212

Annu-Essuman, K. (2014). An Analysis on the Regulation of Grey Market Cyber Materials. *Cornell Internation Affairs Review*, *8*(1). https://doi.org/10.37513/ciar.v8i1.462

Apple security bounty categories. (n.d.). Apple Security Research. https://security.apple.com/bounty/categories

Arora, A., & Telang, R. (2005). Economics of software vulnerability disclosure. *IEEE Security and Privacy Magazine*, *3*(1), 20–25. https://doi.org/10.1109/MSP.2005.12

Bali, A. S., Howlett, M., Lewis, J. M., & Ramesh, M. (2021). Procedural policy tools in theory and practice. *Policy and Society*, *40*(3), 295–311. https://doi.org/10.1080/14494035.2021.1965379

Bardach, E., & Patashnik, E. M. (2020). A practical guide for policy analysis: The eightfold path to more effective problem solving. In *CQ Press; SAGE Publications* (Sixth). SAGE Publications.

Bienz, C., & Juranek, S. (2020). *Software vulnerabilities and bug bounty rograms*. SSRN. https://doi.org/10.2139/ssrn.3599013

Bilge, L., & Dumitraş, T. (2012). Before we knew it: An empirical study of zero-day attacks in the real world. *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, 833–844. https://doi.org/10.1145/2382196.2382284

Böhme, R. (2006). A comparison of market approaches to software vulnerability disclosure. In G. Müller (Ed.), *Emerging Trends in Information and Communication Security* (Vol. 3995, pp. 298–311). Springer Berlin Heidelberg. https://doi.org/10.1007/11766155_21

Brans, M., & Pattyn, V. (2017). Validating methods for comparing public policy: Perspectives from academics and "pracademics". Introduction to the special issue. *Journal of Comparative Policy Analysis: Research and Practice*, *19*(4), 303–312. https://doi.org/10.1080/13876988.2017.1354560

Bugcrowd. (2018). Integrating crowdsourced security with the software development lifecycle. *Bugcrowd*. https://www.bugcrowd.com/blog/integrating-crowdsourced-security-with-the-software-development-lifecycle/

Bugcrowd. (n.d.). *Reporting a bug* [Instruction manual]. Bugcrowd. https://docs.bugcrowd.com/researchers/reporting-managing-submissions/reporting-a-bug/

Çetin, O., Altena, L., Gañán, C., & van Eeten, M. (2018). Let me out! Evaluating the effectiveness of quarantining compromised users in walled gardens. *Proceedings of the Fourteenth Symposium on Usable Privacy and Security*. SOUPS 2018. www.usenix.org/conference/soups2018/presentation/cetin

Choi, J. P., Fershtman, C., & Gandal, N. (2010). Network security: Vulnerabilities and disclosure policy. *The Journal of Industrial Economics*, *58*(4), 868–894. https://doi.org/10.1111/j.1467-6451.2010.00435.x

Cyber attacks statistics. (n.d.). Hackmageddon. https://www.hackmageddon.com/category/security/cy

ber-attacks-statistics/

Cybersecurity and Infrastructure Security Agency. (2013). *National infrastructure protection plan 2013: Partnering for critical infrastructur security and resilience* [Policy plan]. https://www.cisa.gov/resources-tools/resources/2013-national-infrastructure-protection-plan

Cybersecurity and Infrastructure Security Agency. (2020). *BOD 20-01: Develop and publish a vulnerability disclosure policy* (Directive 20–01). https://www.cisa.gov/news-events/directives/bod-20-01-develop-and-publish-vulnerability-disclosure-policy

Directive 2022/2555. (2022). *Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive)*. European Parliament and Council. https://eur-lex.europa.eu/eli/dir/2022/2555

Engin, C. (2023, October 19). Penetration testing vs bug bounty: Compared and explained. *Bugbounter*. https://bugbounter.com/penetration-testing-vs-bug-bounty-compared-and-explained/

Epstein, D., Katzenbach, C., & Musiani, F. (2016). Doing internet governance: Practices, controversies, infrastructures, and institutions. *Internet Policy Review*, *5*(3). https://doi.org/10.14763/2016.3.435

European Commission. (2022). *Cyber Resilience Act* [Summary]. https://digital-strategy.ec.europa.eu/en/library/cyber-resilience-act

European Parliamentary Research Service. (2023). *Investigation of the use of Pegasus and equivalent surveillance spyware* [Plenary report]. https://www.europarl.europa.eu/thinktank/en/document/EPRS_ATA(2023)747923

European Union Agency for Cybersecurity (ENISA). (2022). *Coordinated vulnerability disclosure policies in the EU* [Report]. https://www.enisa.europa.eu/publications/coordinated-vulnerability-disclosure-policies-in-the-eu

European Union Agency for Cybersecurity (ENISA). (2023). *Developing national vulnerabilities programmes* [Report]. https://www.enisa.europa.eu/publications/developing-national-vulnerabilities-programmes

Fadia, A., Nayfeh, M., & Noble, J. (2020). *The role of government cybersecurity efforts in combating risks* [White paper]. McKinsey and Company. https://www.mckinsey.com/industries/public-sector/our-insights/follow-the-leaders-how-governments-can-combat-intensifying-cybersecurity-risks

Fischer, M., & Maggetti, M. (2017). Qualitative comparative analysis and the study of policy processes. *Journal of Comparative Policy Analysis: Research and Practice*, *19*(4), 345–361. https://doi.org/10.1080/13876988.2016.1149281

Flew, T., & Martin, F. R. (Eds.). (2022). *Digital platform regulation: Global perspectives on internet governance*. Palgrave Macmillan. https://doi.org/10.1007/978-3-030-95220-4

Fuster Morell, M. (2022). The gender of the platform economy. *Internet Policy Review*, *11*(1). https://doi.org/10.14763/2022.1.1620

Goldstein, E. (2021, July 29). CISA announces new vulnerability disclosure policy (VDP) platform. *Cybersecurity and Infrastructure Security Agency Blog*. https://www.cisa.gov/news-events/news/cisa-announces-new-vulnerability-disclosure-policy-vdp-platform

Hardin, R., & Cullity, G. (2020). The free rider problem. In E. N. Zalta (Ed.), *The Stanford Encyclopedia*

of Philosophy (Winter 2020 Edition). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/entries/free-rider/

IEEE-CS/ACM Joint Task Force on Software Engineering Ethics and Professional Practices. (n.d.). *Code of ethics* [Code]. Institute of Electrical and Electronics Engineers. https://www.computer.org/education/code-of-ethics

James, A. (2021). The gig economy: A critical introduction. *Economic Geography*, *97*(1), 113–114. https://doi.org/10.1080/00130095.2020.1831908

Jo, A.-M. (2020). Hackers' self-selection in crowdsourced bug bounty programs. *Revue d'économie Industrielle*, *172*, 83–132. https://doi.org/10.4000/rei.9519

Jones, R. L., & Rastogi, A. (2004). Secure coding: Building security into the software development life cycle. *Information Systems Security*, *13*(5), 29–39. https://doi.org/10.1201/1086/44797.13.5.20041101/84907.5

Kestelyn, J. & Bugcrowd Head of Product Marketing. (2022, July 7). How CrowdMatch strengthens crowd engagement and improves researcher rewards. *Bugcrowd*. https://www.bugcrowd.com/blog/how-crowdmatch-strengthens-crowd-engagement-and-improves-researcher-rewards/

Libicki, M. (2015). *Sharing information about threats is not a cybersecurity panacea* (RAND Testimony Series) [Report]. RAND Corporation. https://doi.org/10.7249/CT425

Libicki, M. C., Ablon, L., & Webb, T. (2015). *The defender's dilemma: Charting a course toward cybersecurity*. RAND Corporation. http://www.jstor.org/stable/10.7249/j.ctt15r3x78

Lindblom, C. E. (1959). The science of 'muddling through'. *Public Administration Review*, *19*(2), 79–88. https://doi.org/10.2307/973677

Lindblom, C. E. (1979). Still muddling, not yet through. *Public Administration Review*, *39*(6), 517–526. https://doi.org/10.2307/976178

Maillart, T., Zhao, M., Grossklags, J., & Chuang, J. (2017). Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. *Journal of Cybersecurity*, *3*(2), 81–90. https://doi.org/10.1093/cybsec/tyx008

Malladi, S. S., & Subramanian, H. C. (2020). Bug bounty programs for cybersecurity: Practices, issues, and recommendations. *IEEE Software*, *37*(1), 31–39. https://doi.org/10.1109/MS.2018.2880508

MalwareTech. (2017, May 13). Finding the kill switch to stop the spread of ransomware. *National Cyber Security Centre Blog*. https://www.ncsc.gov.uk/blog-post/finding-kill-switch-stop-spread-ransomware-0

Martin-Vegue, T. (2021). *Optimizing risk response* [White paper]. Information Systems Audit and Control Association (ISACA). https://store.isaca.org/s/store#/store/browse/detail/a2S4w000004KtiREAS

Maurer, S. M. (2017, March 14). A market-based approach to cyber defense: Buying zero-day vulnerabilities. *Bulletin of the Atomic Scientists*. https://thebulletin.org/2017/03/a-market-based-approach-to-cyber-defense-buying-zero-day-vulnerabilities/

McCracken, G. (2019, April 23). Setting up your program reward ranges. *Bugcrowd*. https://www.bugcrowd.com/blog/setting-up-your-program-reward-ranges/

McGraw, G. (2004). Software security. *IEEE Security & Privacy Magazine*, *2*(2), 80–83. https://doi.org/1

0.1109/MSECP.2004.1281254

McKinney, D. (2007). Vulnerability bazaar. *IEEE Security & Privacy Magazine*, *5*(6), 69–73. https://doi.org/10.1109/MSP.2007.180

Meakins, J. (2019). A zero-sum game: The zero-day market in 2018. *Journal of Cyber Policy*, *4*(1), 60–71. https://doi.org/10.1080/23738871.2018.1546883

Mermoud, A., Keupp, M. M., Huguenin, K., Palmié, M., & Percia David, D. (2019). To share or not to share: A behavioral perspective on human participation in security information sharing. *Journal of Cybersecurity*, *5*(1), Article tyz006. https://doi.org/10.1093/cybsec/tyz006

Microsoft. (n.d.). *Microsoft bounty legal safe harbor*. https://www.microsoft.com/en-us/msrc/bounty-safe-harbor

Miller, C. (2007). The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. *Ixth Workshop on the Economics of Information Security*, *1–10*. https://api.semanticscholar.org/CorpusID:12423218

Ministry of Defence. (2020). *Report a vulnerability on an MOD system* [Guidance]. UK Government. https://www.gov.uk/guidance/report-a-vulnerability-on-an-mod-system

National Institute of Standards and Technology. (2018). *Framework for improving critical infrastructure cybersecurity* (Version 1.1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.04162018

National Institute of Standards and Technology - Computer Security Resource Center. (n.d.). Software vulnerability. In *Glossary*. https://csrc.nist.gov/glossary/term/Software_Vulnerability

National Research Council. (2014). *At the nexus of cybersecurity and public policy: Some basic concepts and issues* (p. 18749). National Academies Press. https://doi.org/10.17226/18749

Ohm, M., Plate, H., Sykosch, A., & Meier, M. (2020). Backstabber's knife collection: A review of open source software supply chain attacks. In C. Maurice, L. Bilge, G. Stringhini, & N. Neves (Eds.), *DIMVA 2020: Detection of Intrusions and Malware, and Vulnerability Assessment* (Vol. 12223, pp. 23–43). Springer. https://doi.org/10.1007/978-3-030-52683-2_2

PCI Security Standards Council. (2008). *Payment card industry security standards* [Guidance]. https://www.pcisecuritystandards.org/document_library/?document=pcissc_overview

Peisert, S., Schneier, B., Okhravi, H., Massacci, F., Benzel, T., Landwehr, C., Mannan, M., Mirkovic, J., Prakash, A., & Michael, J. B. (2021). Perspectives on the SolarWinds incident. *IEEE Security & Privacy*, *19*(2), 7–13. https://doi.org/10.1109/MSEC.2021.3051235

Perlroth, N. (2021, February 14). The untold history of America's zero-day market. *Wired*. https://www.wired.com/story/untold-history-americas-zero-day-market/

Poppo, L., Zhou, K. Z., & Li, J. J. (2016). When can you trust "trust"? Calculative trust, relational trust, and supplier performance. *Strategic Management Journal*, *37*(4), 724–741. https://doi.org/10.1002/smj.2374

Powell, W. W., & Snellman, K. (2004). The knowledge economy. *Annual Review of Sociology*, *30*, 199–220. https://doi.org/10.1146/annurev.soc.29.010202.100037

Rajasooriya, S. M., Tsokos, C. P., & Kaluarachchi, P. K. (2016). Stochastic modelling of vulnerability life cycle and security risk evaluation. *Journal of Information Security*, *7*(4), 269–279. https://doi.org/

10.4236/jis.2016.74022

Ramirez, A., Aiello, A., & Lincke, S. J. (2020). A survey and comparison of secure software development standards. *2020 13th CMI Conference on Cybersecurity and Privacy (CMI) - Digital Transformation - Potentials and Challenges(51275)*, 1–6. https://doi.org/10.1109/CMI51275.2020.9322704

Ransbotham, S., Mitra, S., & Ramsey, J. (2012). Are markets for vulnerabilities effective? *MIS Quarterly*, *36*(1), 43–64. https://doi.org/10.2307/41410405

Raymond, E. (1999). The cathedral and the bazaar. *Knowledge, Technology & Policy*, *12*(3), 23–49. https://doi.org/10.1007/s12130-999-1026-0

Schryen, G., & Kadura, R. (2009). Open source vs. closed source software: Towards measuring security. *Proceedings of the 2009 ACM Symposium on Applied Computing*, 2016–2023. https://doi.org/10.1145/1529282.1529731

Schwartz, A., & Knake, R. (2016). *Government's role in vulnerability disclosure: Creating a permanent and accountable vulnerability equities process* (Discussion Paper 2016–03; The Cyber Security Project, p. 28). Harvard Kennedy School - Belfer Center. https://www.belfercenter.org/sites/default/files/legacy/files/vulnerability-disclosure-web-final3.pdf

Shahzad, M., Shafiq, M. Z., & Liu, A. X. (2012). *A large scale exploratory analysis of software vulnerability life cycles*. 771–781. https://doi.org/10.1109/ICSE.2012.6227141

Soleimani, M. (2022). Buyers' trust and mistrust in e-commerce platforms: A synthesizing literature review. *Information Systems and E-Business Management*, *20*(1), 57–78. https://doi.org/10.1007/s10257-021-00545-0

Souppaya, M., Scarfone, K., & Dodson, D. (2022). *Secure software development framework (SSDF) version 1.1: Recommendations for mitigating the risk of software vulnerabilities* (Report 800–218; NIST Special Publication, p. NIST SP 800-218). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-218

Subramanian, H. C., & Malladi, S. (2020). Bug bounty marketplaces and enabling responsible vulnerability disclosure: An empirical analysis. *Journal of Database Management*, *31*(1), 38–63. https://doi.org/10.4018/JDM.2020010103

Terwiesch, C., & Xu, Y. (2008). Innovation contests, open innovation, and multiagent problem solving. *Management Science*, *54*(9), 1529–1543. https://doi.org/10.1287/mnsc.1080.0884

The disclose.io Project. (n.d.). https://disclose.io/

The Electronic Privacy Information Center. (n.d.). *Vulnerabilities equities process*. https://archive.epic.org/privacy/cybersecurity/vep/

The National Cyber Security Centre. (2020). *Vulnerability disclosure toolkit* [Toolkit]. UK Government. https://www.ncsc.gov.uk/information/vulnerability-disclosure-toolkit

The White House. (2017). *Vulnerabilities equities policy and process for the United States Government* [Policy report]. https://www.hsdl.org/c/abstract/?docid=805726

The White House. (2023). *National cybersecurity stratergy* [Strategy]. https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf

United States-Computer Emergency Readiness Team. (n.d.). [Document]. Homeland Security. https://w

ww.cisa.gov/sites/default/files/publications/infosheet_US-CERT_v2.pdf

Varian, H. (2004). System reliability and free riding. In L. J. Camp & S. Lewis (Eds.), *Economics of Information Security* (Vol. 12, pp. 1–15). Kluwer Academic Publishers. https://doi.org/10.1007/1-4020-8090-5_1

Wachs, J. (2022). *Making markets for information security: The role of online platforms in bug bounty programs* (arXiv:2204.06905; Version 1). arXiv. https://doi.org/10.48550/ARXIV.2204.06905

Walshe, T., & Simpson, A. (2020). *An empirical study of bug bounty programs*. 35–44. https://doi.org/10.1109/IBF50092.2020.9034828

Walshe, T., & Simpson, A. (2023). Towards a greater understanding of coordinated vulnerability disclosure policy documents. *Digital Threats: Research and Practice*, *4*(2), 1–36. https://doi.org/10.1145/3586180

Walshe, T., & Simpson, A. C. (2022). Coordinated vulnerability disclosure programme effectiveness: Issues and recommendations. *Computers & Security*, *123*, Article 102936. https://doi.org/10.1016/j.cose.2022.102936

Weimer, D. L., & Vining, A. R. (2017). *Policy analysis: Concepts and practice* (6th ed.). Routledge. https://doi.org/10.4324/9781315442129

Williamson, O. E. (2008). Outsourcing: Transaction cost economics and supply chain management. *Journal of Supply Chain Management*, *44*(2), 5–16. https://doi.org/10.1111/j.1745-493X.2008.00051.x

World Economic Forum. (2023). Digital dependencies and cyber vulnerabilities. In *Global risks report 2023* (18th ed.). World Economic Forum. https://www.weforum.org/reports/global-risks-report-2023/in-full/chapter-3-digital-dependencies-and-cyber-vulnerabilities

Yadav, D. (2022). Criteria for good qualitative research: A comprehensive review. *The Asia-Pacific Education Researcher*, *31*, 679–689. https://doi.org/10.1007/s40299-021-00619-0

Zanini, M. T. F., & Musante, M. (2013). Trust in the knowledge economy. *Journal of Business & Industrial Marketing*, *28*(6), 487–493. https://doi.org/10.1108/JBIM-04-2013-0102

Zenger, T. R., Felin, T., & Bigelow, L. (2011). Theories of the firm–market boundary. *Academy of Management Annals*, *5*(1), 89–133. https://doi.org/10.5465/19416520.2011.590301

Zhao, M., Laszka, A., & Grossklags, J. (2017). Devising effective policies for bug-bounty platforms and security vulnerability discovery. *Journal of Information Policy*, *7*, 372–418. https://doi.org/10.5325/jinfopoli.7.2017.0372

Zrahia, A. (2018). Threat intelligence sharing between cybersecurity vendors: Network, dyadic, and agent views. *Journal of Cybersecurity*, *4*(1). https://doi.org/10.1093/cybsec/tyy008

Zrahia, A., Gandal, N., Markovich, S., & Riordan, M. H. (2022). *The simple economics of an external dhock on a crowdsourced 'bug bounty platform'*. SSRN. https://doi.org/10.2139/ssrn.4154516

Published by
Alexander von Humboldt Institute for Internet and Society

in cooperation with
CREATe
cnrs centre internet et societe
R&I IN3 Internet interdisciplinary Institute Universitat Oberta de Catalunya
University of Tartu Johan Skytte Institute of Political Studies