Economics of vulnerability disclosure

DECEMBER 2018

www.enisa.europa.eu
European Union Agency For Network and Information Security

About ENISA

The European Union Agency for Network and Information Security (ENISA) is a centre of network and information security expertise for the EU, its member states, the private sector and EU citizens. ENISA works with these groups to develop advice and recommendations on good practice in information security. It assists member states in implementing relevant EU legislation and works to improve the resilience of Europe’s critical information infrastructure and networks. ENISA seeks to enhance existing expertise in member states by supporting the development of cross-border communities committed to improving network and information security throughout the EU. More information about ENISA and its work can be found at www.enisa.europa.eu.

Authors

This report was written by Erik Silfversten, William Phillips, Giacomo Persi Paoli (RAND Europe) and Cosmin Ciobanu (ENISA).

Contact

For queries in relation to this paper, please use opsec@enisa.europa.eu
For media enquiries about this paper, please use press@enisa.europa.eu.

Acknowledgements

The authors of this report are particularly thankful for the time and valuable insights given by the interviewees, which have significantly contributed to achieving a more well-informed and balanced study. The full list of interviewees can be found in Annex A of this document.

Legal notice

Notice must be taken that this publication represents the views and interpretations of ENISA, unless stated otherwise. This publication should not be construed to be a legal action of ENISA or the ENISA bodies unless adopted pursuant to the Regulation (EU) No 526/2013. This publication does not necessarily represent state-of the-art and ENISA may update it from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for information purposes only. It must be accessible free of charge. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

Copyright Notice

ISBN: 978-92-9204-277-6, DOI: 10.2824/49807

Table of Contents

Executive Summary 5

1. Introduction 7

1.1 Background to the study 7

1.2 Objectives and structure of the study 7

1.3 Caveats and limitations 8

1.4 Outline of the report 8

2. Overview of vulnerability disclosure 9

2.1 Definitions and key concepts 9

2.2 Actors in vulnerability disclosure 10

2.3 Types of vulnerability disclosure processes 12

2.3.1 It is important to understand the differences between various vulnerability disclosure mechanisms and how they interact with other information security interventions 14

2.4 The nature of vulnerabilities 15

3. Introduction to information security economics 18

3.1 The information security market 18

3.1.1 Network effects mean that products and services are subject to different security considerations in relation to their size and use 19

3.1.2 Externalities may have significant consequences for behaviour in the information security market 19

3.1.3 Liability dumping may enable actors to shift responsibilities and costs to other actors 20

3.1.4 A lack of or variability of knowledge results in sub-optimal choices in the information security market 20

3.1.5 Moral hazard means that decision makers in the information security market take decisions subject to high risks 21

3.1.6 The economic features of the information security market affect the security of products and services, but also influence actors within the vulnerability disclosure landscape 22

3.2 The changing nature of the information security market 22

4. Incentives and behaviour in vulnerability disclosure 26

4.1 Introduction 26

4.2 The nature of incentives 26

4.3 Finder behaviour 27

4.3.1 Internal incentives act as the primary driver for finder behaviour in the vulnerability disclosure process 27

4.3.2 Finders also respond to a number of external economic considerations and incentives 30

4.4 Organisational considerations and incentives 33

4.4.1 Private and public sector organisations respond to regulatory and user expectations while operating with limited resources 33

4.4.2 Governments may take on several different roles in the vulnerability disclosure process, which affects its incentives and considerations 39

4.5 Structural considerations and incentives 41

4.5.1 Legislation plays a central role in the vulnerability disclosure landscape 41

4.5.2 Regulation and liability can provide structure to the vulnerability disclosure process and incentivise desired behaviours 42

4.5.3 Insurance represents another type of structural incentive that can influence behaviour 44

4.6 Normative considerations and incentives 45

4.7 Emerging trends in vulnerability disclosure 46

4.8 On reflection 49

5. Vulnerability case studies 51

5.1 Meltdown and Spectre 51

5.1.1 Meltdown and Spectre were discovered by three security teams working independently of each other 51

5.1.2 The disclosure process highlights the complexity in multivendor CVD 52

5.1.3 Meltdown and Spectre have had an extensive impact on manufacturers due to remediation costs rather than attack costs 53

5.1.4 The discovery of Meltdown and Spectre highlights the importance of investing in long-term security research 54

5.2 EternalBlue 55

5.2.1 Discovery and disclosure of EternalBlue represents a non-disclosure process 55

5.2.2 EternalBlue had significant and wide-ranging global implications 56

5.2.3 The EternalBlue disclosure highlights the profound importance of responsible disclosure 57

5.3 On reflection 59

6. Summary and key findings 61

6.1 Summary 61

6.2 Key findings 61

7. References 64

Annex A: List of interviewees 73

Executive Summary

Vulnerability disclosure refers to the process of identifying, reporting and patching weaknesses of software, hardware or services that can be exploited. The different actors within a vulnerability disclosure process are subject to a range of economic considerations and incentives that may influence their behaviour. These economic aspects of vulnerability disclosure are often overlooked and poorly understood, but may help explain why some vulnerabilities are disclosed responsibly while others are not. This study serves as a follow up to the 2015 ENISA Good Practice Guide on Vulnerability Disclosure[1] and has the overarching objective to improve the understanding of the economics of vulnerability disclosure by providing a glimpse into the costs, incentives and impact related to discovering and disclosing vulnerabilities. This objective was pursued through a mixed-methods approach, which included desk research, literature review, case studies and expert interviews.

Vulnerability disclosure takes place in a wider computing and information security ecosystem whose unique economic structures and incentives have direct economic effects on vulnerability disclosure. This extends to both the structure and parameters of the market for information security, as well as the nature of software and hardware vulnerabilities. This wider ecosystem is also subject to continuous change driven by technological development to which vulnerability disclosure must adapt and respond.

Economic decisions taken in the vulnerability disclosure process largely depend on the particular incentives perceived by each actor at different stages of the process. There are four main actor groups within the vulnerability disclosure process: Users, Vendors, Finders and Coordinators. There are also several possible vulnerability disclosure options that actors can engage in, including full, limited or non-disclosure, which further influence the types of economic considerations and incentives that are present. Vulnerability disclosure actors are subject to economic incentives and motivations that may influence their behaviour at the individual and organisational level, as well as at the structural and normative levels. It is clear that economic incentives play a key role in vulnerability disclosure across all actors and processes, regardless of what type of vulnerability disclosure process is ultimately pursued.

Many of the behaviours and incentives in vulnerability disclosure are also affected by both negative and positive external factors, which highlights that behaviour in vulnerability disclosure can be influenced and changed through different mechanisms. Structural levers, such as legislation or regulation, can be important policy tools to influence the behaviour of different vulnerability disclosure actors to achieve socially desirable security outcomes. Legislation and regulation may help offset some of the negative consequences of the economic features of the information security market.

The two case studies of Meltdown/Spectre and EternalBlue featured in this report highlight the diversity found in vulnerability disclosure and illustrate the distinct, and potentially grave, differences between a coordinated vulnerability disclosure and a non-disclosure process. While both case studies showed the costs that vulnerabilities can incur, they also illustrated the cost savings that responsible disclosure can realise by reducing the exploitation of identified vulnerabilities.

The prevalence of economics in vulnerability disclosure emphasises the importance of a well-developed understanding of the economic aspects of vulnerability disclosure and how these aspects may influence different processes and actor behaviour within them. However, currently there are research gaps in

[1] See https://www.enisa.europa.eu/publications/vulnerability-disclosure (As of 31 October 2018).

certain areas of the research field that may impede a full understanding of the economics of vulnerability disclosure.

Overall, the study has produced a number of key findings and recommendations:

- First and foremost, the study shows the importance that vulnerability disclosure, particularly coordinated vulnerability disclosure, plays in modern society. Vulnerabilities in widely-used software and hardware can cause immense societal harm and it is necessary to have processes in place to adequately identify, report, receive, triage and mitigate vulnerabilities.

- All actors involved in vulnerability disclosure should therefore recognise the importance of setting up and running appropriate and mutually beneficial structures that enables effective and efficient coordinated vulnerability disclosure to take place.

- Industry that develop or manufacture products or services for the Internet or the global ICT ecosystem should seek to ensure the ability to receive good-faith vulnerability reports from the community and, as appropriate, mitigate them.

- National governments should consider implementing a coordinated vulnerability disclosure policy and if appropriate, embark on a discussion of how to best approach a government disclosure decisions process.

- Awareness raising and capacity building across all actor groups are key enablers for a well-functioning vulnerability disclosure ecosystem and for actors to understand the economic incentives and behaviour of other parties involved. Providing actors with resources, good practice and voluntary standards are important tools to consider in promoting mutually beneficial and standardised behaviour.

- There are also opportunities to improve finder wellbeing and the overall vulnerability disclosure ecosystem by ensuring safe harbour practices and legal safeguards for security researchers working in good-faith to identify and report vulnerabilities.

- Lastly, the study recognises that vulnerability disclosure is one part of a larger information security ecosystem and encourages continuous efforts to improve the quality and security of software and hardware to reduce the number of vulnerabilities in deployment, as well as continuous investment in long-term security research to identify and mitigate fundamental weaknesses such as design flaws or protocol vulnerabilities.

The analysis presented in this report will be useful to all the key stakeholders involved or affected to some extent by a vulnerability disclosure in a software, hardware component or system, including researchers, consumers, vendors, vulnerability coordinators and brokers, regulators, managers, information security experts and officers.

1. Introduction

1.1 Background to the study

In network and information security, a vulnerability is defined as a weakness of software, hardware or a service that can be exploited. In recent years, there have been numerous high-profile vulnerabilities disclosed or exploited that have incurred significant economic and societal costs. In 2018, exploitation of the EternalBlue vulnerability through the WannaCry and Petya/NotPetya ransomware attacks resulted in significant societal disruption in Europe and beyond. The same year also saw the disclosure of the Spectre and Meltdown vulnerabilities, which affected nearly all computer chips manufactured in modern history.[2]

While some vulnerabilities have been responsibly disclosed following due process, other vulnerabilities have come to light only after substantial associated cyber attacks have resulted. This has led the information security community to question whether the manner of disclosure was appropriate, especially since there often appears to be a connection between the public disclosure of the vulnerability and the subsequent levels of exploitation. The different actors within a vulnerability disclosure process are subject to a range of economic considerations and incentives that may influence their behaviour. These economic aspects of vulnerability disclosure are often overlooked and poorly understood, but may help explain why some vulnerabilities are disclosed responsibly while others are not.

In the context of ENISA's strategic objective to develop and maintain a high level of expertise of EU actors taking into account evolutions in Network & Information Security (NIS),[3] ENISA is following up on the 2015 Good Practice Guide on Vulnerability Disclosure[4] to better understand the economics of vulnerability disclosure.

1.2 Objectives and structure of the study

The objective of this study was to provide a glimpse into the costs, incentives and impact related to discovering and disclosing vulnerabilities. The main audience for this document is comprised of all the key stakeholders involved or affected to some extent by a vulnerability disclosure in a software, hardware component or system, including researchers, consumers, vendors, vulnerability coordinators and brokers, regulators, managers, information security experts and officers.

The study was carried out through a mixed-methods approach comprising desk research, literature review and key informant interviews. A review of the available literature – including academic research, technical reports, company publications, media articles, and blogs – provided the basis for the discussion in Chapters 2–4. In parallel, a set of case studies was prepared that focused on two widely-distributed and critical vulnerabilities disclosed in 2017.

In addition, the research team also conducted a total of 13 interviews with experts from the vulnerability disclosure community, including representatives from academia, bug bounty platforms,[5] vulnerability disclosure programme operators, vendors and others. A full list of the study interviewees can be found in Annex A.

[2] See Chapter 5 for additional information on EternalBlue, Meltdown and Spectre.
[3] ENISA (2018).
[4] ENISA (2015).
[5] Platforms in this context refers to software or online services used to deploy and operate bug bounty programmes.

1.3 Caveats and limitations

This study is subject to a number of caveats and limitations:

- The economics of vulnerability disclosure is an emerging research field, in which some aspects have been subject to more enquiry than others. As such, there are sections of this report that have been informed by a comparatively limited evidence base. In these instances, additional interview data was gathered to complement the desk research and literature review results. The study team has also made a conscious effort to note particular areas where additional research could be beneficial to European stakeholders.
- The study team also gathered additional perspectives on the economics of vulnerability disclosure through interviews with vulnerability disclosure stakeholders. While these interviewees represent a range of different stakeholders, it is important to bear in mind that these interviews do not represent a complete view of all stakeholders and individuals involved in vulnerability disclosure.

1.4 Outline of the report

Beyond this introductory chapter, the report comprises five further chapters:

2. Overview of vulnerability disclosure features a brief introduction to vulnerability disclosure and presents key definitions, processes and actors.

3. Introduction to the economics of vulnerability disclosure examines the economic aspects of the information security market and how this relates to vulnerability disclosure. This chapter also features a discussion of classical economics concepts that have clear relevance and application to the issue of vulnerability disclosure.

4. Incentives and behaviour in vulnerability disclosure explores how different incentives affect the behaviour of actors in vulnerability disclosure and how other external factors may also affect the wider vulnerability disclosure ecosystem.

5. Vulnerability case studies features two case studies of recently disclosed high-profile, vulnerabilities. The chapter highlights how vulnerability disclosure may happen in practice and what role economic considerations may play.

6. Summary and key findings comprises a short summary of the study’s key findings.

This report is also accompanied by Annex A, which features a list of experts interviewed for the study.

2. Overview of vulnerability disclosure

Software, hardware and online services are all susceptible to vulnerabilities and it is unlikely that vulnerabilities will ever be completely eradicated. Even if a system is sufficiently secure at launch, there is no guarantee it will remain that way. Deployment in a new context, interactions with new systems or development of new attack methods may uncover previously unknown vulnerabilities.[6]

Vulnerabilities can be caused by a number of factors, including design and development flaws,[7] misconfiguration, inadequate administrative or operational processes, other user errors, or unforeseen changes in the operating environment or threat landscape. In an interconnected environment or network, the presence of vulnerabilities in popular software, hardware or services may present considerable risk to systems and society – which calls for efficient identification and remediation of vulnerabilities. Vulnerabilities that go undetected for a prolonged period of time or are inappropriately disclosed[8] may further exacerbate these risks, prompting the need for effective vulnerability disclosure processes.

This chapter briefly explains the key concepts, actors and processes involved in vulnerability disclosure. For a more comprehensive overview of the vulnerability disclosure landscape, challenges and good practice please see the associated ENISA publication on Good Practice to Vulnerability Disclosure.[9]

2.1 Definitions and key concepts

There are two ISO standards related to vulnerability management: ISO/IEC 29147: Vulnerability disclosure, and ISO/IEC 30111: Vulnerability handling processes. These define key concepts and provide guidelines to vendors for processes related to the receipt and handling of vulnerability information. Within the context of this study, the following definitions outlined in Table 2.1 have been used.[10]

Table 2.1 Definitions of key terms

Term	Definition
Advisory	An announcement or bulletin that serves to inform, advise and warn about a vulnerability of a product or service.
Disclosure	The act of initially providing vulnerability information to a party that was not believed to be previously aware. The overall disclosure process typically includes multiple disclosure events.
Remediation	Patch, fix, upgrade, configuration or documentation change to either remove or mitigate a vulnerability, typically provided by vendors.

[6] Interview 12, Householder et al. (2017).

[7] For a comprehensive overview of common potential weaknesses in software, see the community-developed Common Weakness Enumeration (CWE) list. CWE is hosted by the MITRE Corporation and sponsored by the United States Computer Emergency Response Team (US-CERT) in the office of Cybersecurity and Communications at the US Department of Homeland Security.

[8] Scenarios in which a vulnerability is disclosed too quickly, before appropriate remediation steps can be taken, or too slowly, where the vulnerability is left unpatched, may present significant security risk.

[9] Available through the ENISA website: https://www.enisa.europa.eu/publications/vulnerability-disclosure (As of 31 October 2018).

[10] Definitions available in ISO 29147 have been adopted with minimal modifications.

Vulnerability	A weakness of software, hardware or online service that can be exploited.
Zero-day	A vulnerability for which no patch or fix has been publicly released.

Source: ISO 29147 and FIRST (2017).

Economic decisions taken in the vulnerability disclosure process largely depend on the particular incentives perceived by each actor at different stages of the process. Vulnerability disclosure can be conducted through a number of different processes and may involve a number of different actors, as presented in further detail below.

2.2 Actors in vulnerability disclosure

There are four main actor groups within the vulnerability disclosure process:[11]

- **Users** of software, hardware and services, and may refer to individuals, organisations or governments.
- **Vendors** that comprise the developers, manufacturers and suppliers of software, hardware and services. This may also include so-called ‘intermediate vendors’ that make up the supply chain of a specific product or service.
- **Finders** who make up the community of individuals that identify and report vulnerabilities. Finders are sometimes also referred to as discoverers, reporters or researchers.
- **Coordinators** are trusted organisations that act as intermediaries between finders and vendors to ensure that vulnerabilities are disclosed and mitigated responsibly. Well-known coordinators include Computer Emergency Response Teams (CERT) such as the US-CERT Coordination Center (CERT/CC), Finland CERT (CERT-FI) and Japan CERT (JP-CERT).

Additionally, there are also a number of secondary actors who perform indirect roles in the vulnerability disclosure process, including governments, the media and adversarial actors.[12]

- **Governments** play a complex role in the vulnerability disclosure process. They can act as finders, vendors and coordinators, as well as acquire or maintain vulnerabilities for national security purposes. Governments also develop legislation and regulations that may influence vulnerability disclosure.
- **Media** reports on vulnerabilities and engages in the dissemination of vulnerability information.
- **Adversarial actors** such as organised criminals or other adversaries may exploit vulnerabilities or engage in the vulnerability disclosure process for nefarious purposes.

Figure 2.1 Overview of the actors involved in vulnerability disclosure grouped to typical functions

[11] ENISA (2015).

[12] ENISA (2015).

Overview of the actors involved in vulnerability disclosure grouped to typical functions

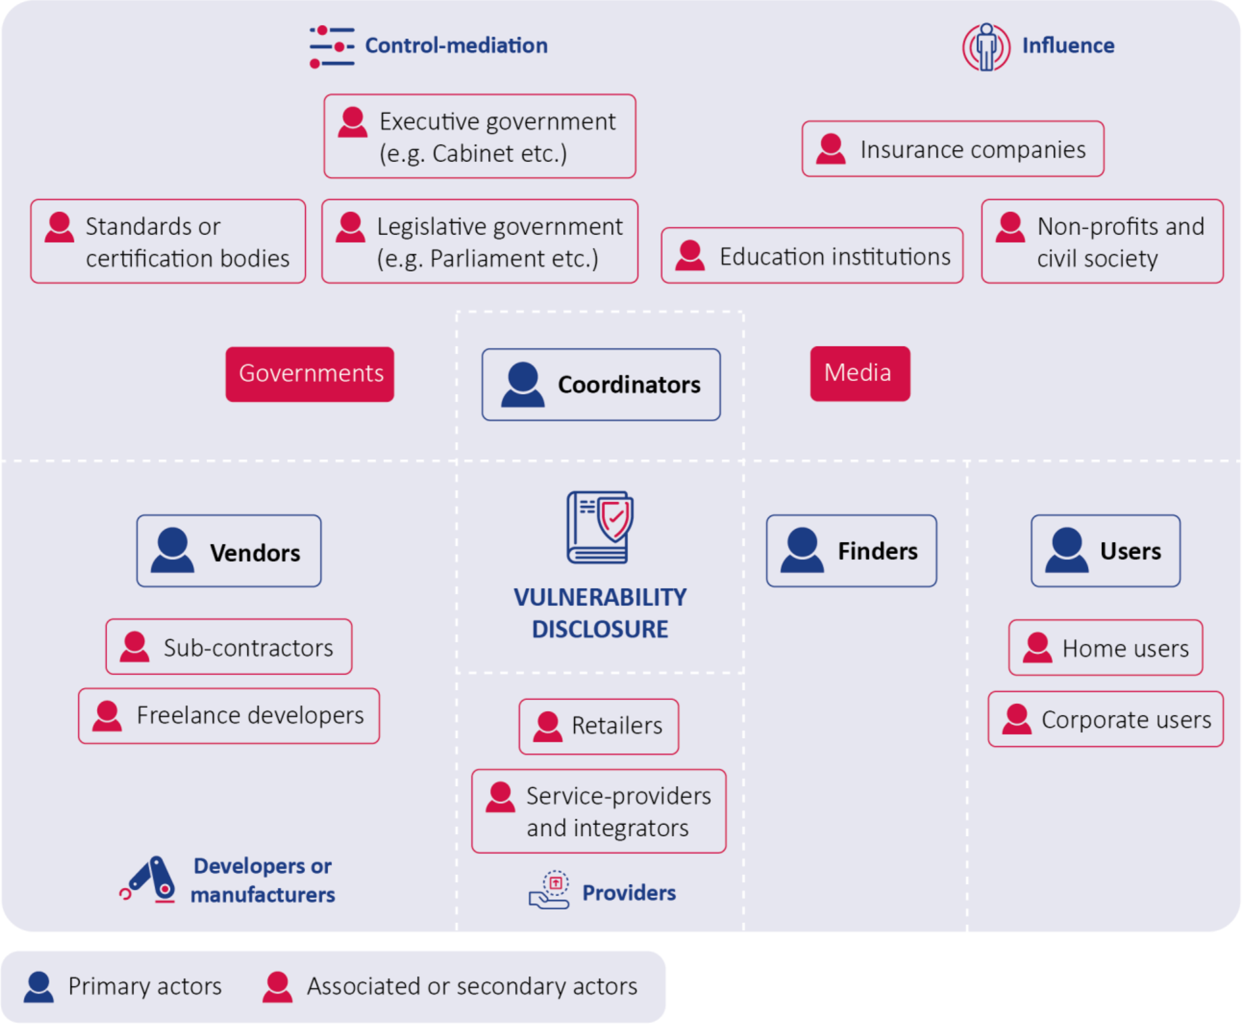
Control-mediation
Influence
Executive government (e.g. Cabinet etc.)
Insurance companies
Standards or certification bodies
Legislative government (e.g. Parliament etc.)
Education institutions
Non-profits and civil society
Governments
Coordinators
Media
Vendors
VULNERABILITY DISCLOSURE
Finders
Users
Sub-contractors
Freelance developers
Home users
Corporate users
Retailers
Service-providers and integrators
Developers or manufacturers
Providers
Primary actors
Associated or secondary actors

At the same time, the exact nature of the steps taken in vulnerability disclosure and the actors involved ultimately depend on the type of vulnerability disclosure process followed. An overview of these is provided below.

2.3 Types of vulnerability disclosure processes

As shown in Figure 2.2, there are several possible vulnerability disclosure options that actors can engage in, which further influence the types of economic considerations and incentives that are present. There are three possible courses of action for disclosing a vulnerability: full, limited or non-disclosure.[13]

- **Full disclosure** refers to when an identifier releases all information about an identified vulnerability publicly, without coordinating with or waiting for coordinator or vendor action.
- **Limited disclosure** refers to when an identifier works with a coordinator or vendor to minimise the risk of the identified vulnerability. Once a patch has been developed, the coordinator or vendor will publish the vulnerability information alongside the remediation measures. Limited disclosure may also be referred to as responsible or coordinated disclosure.
- **Non-disclosure** may occur due to a number of reasons. Individual finders may opt to not disclose vulnerabilities in exchange for payments, particularly if higher payouts can be achieved in the black market compared to responsible disclosure avenues.[14] Another emerging area of non-disclosure is related to government-run initiatives to analyse, evaluate and select vulnerabilities to keep secret for national security purposes, which are sometimes referred to as vulnerability equities processes. The reasoning is that a government may wish not to disclose information about particular vulnerabilities in order to exploit those vulnerabilities for intelligence gathering or for other offensive cyber operations.[15]

Limited disclosure often occurs through coordinated or responsible disclosure. **Coordinated vulnerability disclosure (CVD)** is a process where vulnerability finders work with either vendors or coordinators to minimise the risk of an identified vulnerability and typically involves a set of steps that require careful management so as to avoid potential negative impacts, which otherwise could be substantial.[16] In general, CVD aims to:

1. Ensure that an identified vulnerability is addressed by the vendors.
2. Minimise the risk that the vulnerability presents.
3. Provide users with adequate information to evaluate the risk posed to their systems by the identified vulnerability.
4. Set expectations to encourage positive communication and coordination among the involved actors and stakeholders.[17]

This study deals primarily with CVD and, unless otherwise specified, vulnerability disclosure refers to CVD.

[13] ENISA (2015).
[14] Ablon & Bogart (2017).
[15] ENISA (2015).
[16] Tang et. al. (2016).
[17] Householder et. al. (2017).

Figure 2.2 High-level overview of vulnerability disclosure processes and flows

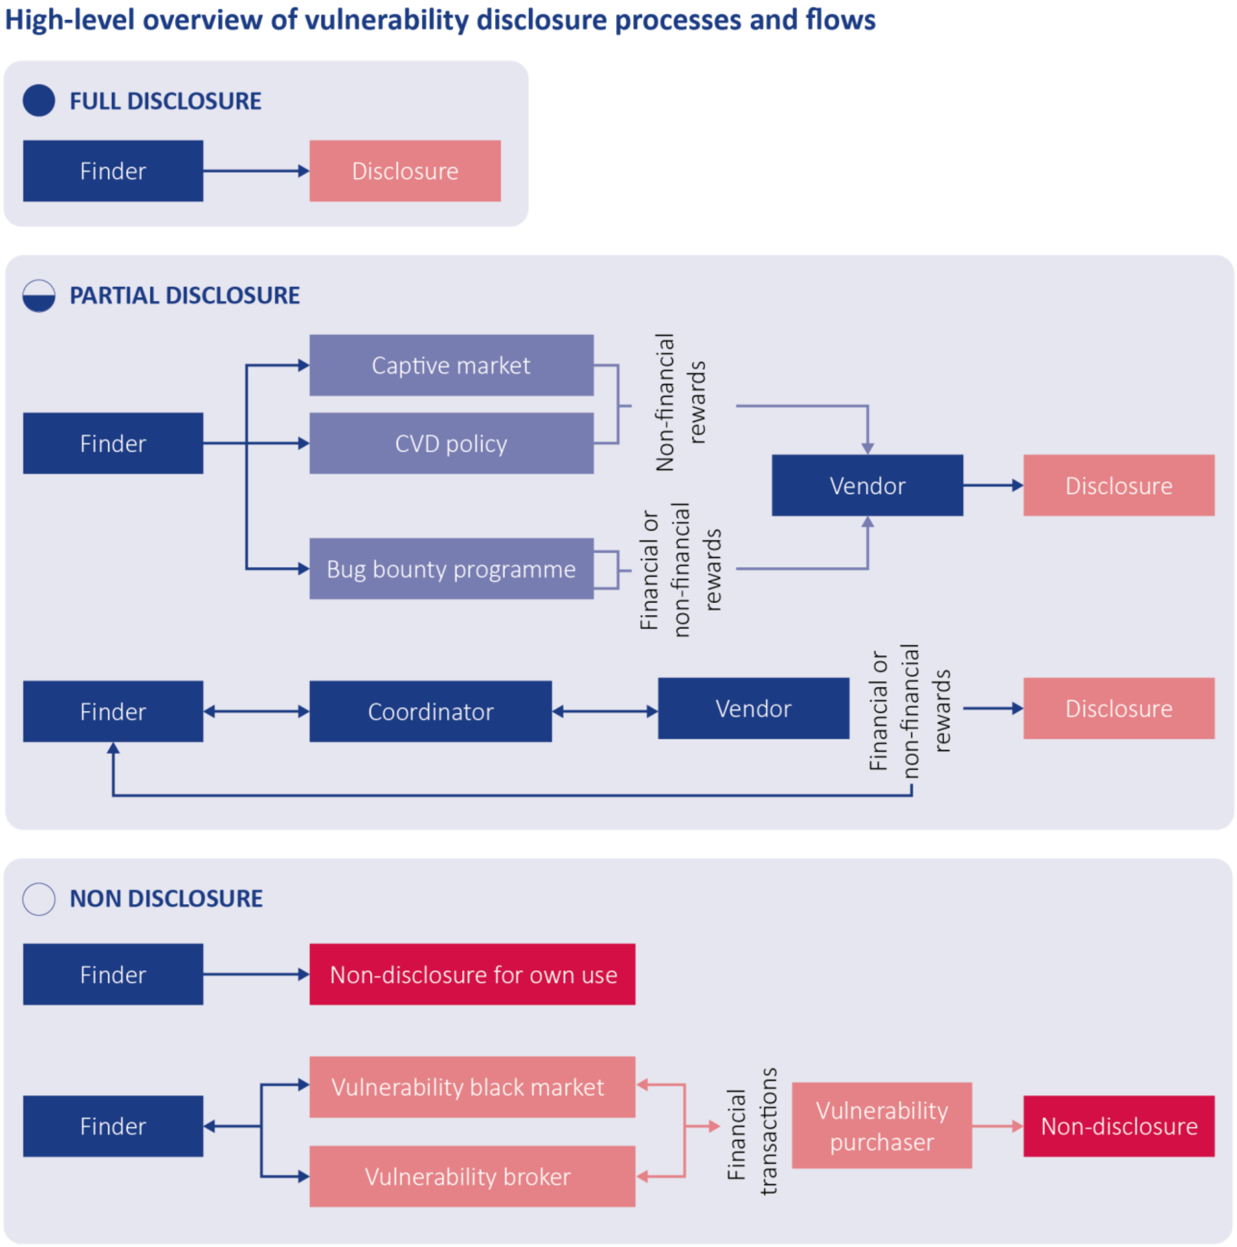

Source: ENISA study on the Economics of Vulnerability Disclosure

Limited or non-disclosure may also take place through so-called 'vulnerability markets'. A market is made up of commodities – in this case undisclosed vulnerabilities, which are sold by their producers (i.e. finders) to consumers (i.e. vendors, governments or malicious actors).[18] A vulnerability market may be unregulated or regulated. An unregulated market is characterised by few rules or limitations, and sales are typically made to the highest bidder. In contrast, regulated markets usually comprise set rules and processes with

[18] Shahzad et al. (2012); Arora & Telang (2005); Algarni & Malaiya (2014).

which sellers need to comply, and that may restrict sales to particular customer groups (e.g. only selling to governments).[19]

Regulated markets can take the following forms:

- **Coordinated disclosure markets**, where vulnerabilities are publicly disclosed either through the vendor or a coordinator (such as a CERT). Finders may or may not receive financial or non-financial rewards for disclosure.
- **Captive markets**, where finders disclose vulnerabilities to the vendor or host organisation and the vulnerability is not disclosed publicly. This may entail security researchers working within or under contract for a particular organisation, as well as security researchers that work for government agencies in defence or intelligence services.
- **Vulnerability rewards markets**, where finders disclose vulnerabilities through a vendor or trusted third party in exchange for financial or non-financial rewards, for example through a bug bounty programme. Rewards are typically linked to the severity of the vulnerability and its potential security implications. Bug bounty programmes may include vendor-specific programmes,[20] bug bounty platforms[21] or coordinated vulnerability rewards programmes.[22]

In addition, unregulated markets include:

- **Partially regulated markets**, such as vulnerability brokers who provide a link between sellers and buyers and typically take a commission when sales are finalised. **Brokers** are organisations or individuals that receive vulnerability information from finders and facilitate in the finding of a buyer for a particular vulnerability. Vulnerability brokers may have certain rules of conduct or limitations, but typically sell vulnerabilities to the highest bidder.[23] Brokers tend to focus their efforts on zero-day vulnerabilities and focus their sales to government agencies.
- **Vulnerability black markets**, which are unregulated markets with certain attributes such as unknown buyers, non-excludability[24] of vulnerabilities in the market, reliance on personal connections to trade, and absences of guarantees to keep the vulnerability secret. These types of markets may occur in different locations such as online chat rooms, marketplace websites or on the dark web.[25]

2.3.1 It is important to understand the differences between various vulnerability disclosure mechanisms and how they interact with other information security interventions

Terms related to vulnerability disclosure are sometimes conflated and used interchangeably, particularly in relation to CVD and bug bounties. A CVD policy is not equivalent or the same as a bug bounty programme.

[19] Miller (2007); Algarni & Malaiya (2014).

[20] Bug bounty programmes run by an individual organisation or vendor such as those operated by Mozilla (https://www.mozilla.org/en-US/security/bug-bounty/), Google (https://bughunter.withgoogle.com/), or Facebook (https://www.facebook.com/whitehat) (As of 31 October 2018).

[21] Such as BugCrowd (https://www.bugcrowd.com/), HackerOne (https://www.hackerone.com/), BountyFactory (https://bountyfactory.io) or Intigriti (https://www.intigriti.com) (As of 31 October 2018).

[22] Such as the Zero-day Initiative (https://www.zerodayinitiative.com/) (As of 31 October 2018).

[23] Such as Zerodium (https://zerodium.com/) (As of 31 October 2018).

[24] In economics, a good is non-excludable if non-paying consumers cannot be prevented from accessing it. In this case it may mean that even though a consumer has purchased a vulnerability on a black market, the seller may sell the same vulnerability again to another customer.

[25] Radianti (2010).

A CVD policy is the primary good practice mechanism that enables security researchers or finders to report identified potential vulnerabilities to an organisation through a dedicated and structured vulnerability reporting channel. In contrast, a bug bounty programme allows organisations to define and scope a programme where security researchers are allowed to try to identify security vulnerabilities – often within a subset of the organisation's technical infrastructure – in exchange for financial or non-financial 'bounties' for successfully validated vulnerabilities.

A bug bounty programme can therefore be a useful tool to incentivise the researcher community to hunt for vulnerabilities as a complement to a CVD policy, but such a programme should not be seen as a replacement for a CVD policy. At the same time, several interviewees emphasised the importance of recognising that CVD is not a replacement for other forms of security measures or interventions.[26] Organisations should engage in CVD in addition to investing in appropriate information security arrangements and other forms of security testing (e.g. secure-by-design, penetration tests, security audits, etc.).

2.4 The nature of vulnerabilities

The economics of vulnerability disclosure are also influenced by the characteristics and features of vulnerabilities. Three categories of vulnerability characteristics that may also influence economic considerations and incentives in the vulnerability disclosure are:

1. The prevalence of vulnerabilities.
2. The rediscovery rates and lifespan of vulnerabilities.
3. The cost and time needed to exploit vulnerabilities.

First, **vulnerabilities are common**. The number of reported vulnerabilities has increased in recent years and several high-profile vulnerabilities have been uncovered, either through responsible disclosure or through their exploitation.[27] It is also likely that there are more vulnerabilities than those that have been reported so far. However, the frequency of vulnerabilities should not be the only metric that is used when considering the severity of the threat posed by the presence of vulnerabilities; not all vulnerabilities are created equal. Simply looking at the number of vulnerabilities does not reveal the potential severity of the vulnerability (i.e. what damage or patching costs it could incur), if the vulnerability has been exploited by a malicious actor, or if the vulnerability has been patched by the vendor.[28]

Vulnerabilities also have **varying rediscovery rates**. The likelihood that two or more security researchers identify a vulnerability independently from each other is referred to as the collision rate (or overlap rate). The threat of another security researcher or malicious actor rediscovering a vulnerability may serve as an incentive for vendors to patch the vulnerability as soon as possible.[29] However, the rate or likelihood of this type of independent rediscovery is contested and may warrant additional research. Collision rates could, in theory, be limited given the complexity of software and the high number of vulnerabilities that are continuously identified.[30]

As shown in Figure 2.3, a study of zero-day vulnerabilities – which are often perceived as being potentially the most harmful – found that collision rates vary significantly over time. This means that the lifespan of

[26] Interviews 4, 8.

[27] See Chapter 5 for two case studies of recent high-profile vulnerabilities.

[28] ENISA (2015).

[29] Arora et al. (2008); Cavusoglu et. al. (2005).

[30] Rescorla (2005).

zero-day vulnerabilities and their associated exploits could be relatively long – up to an average of seven years.[31]

Figure 2.3 Median collision rates for zero-day vulnerabilities

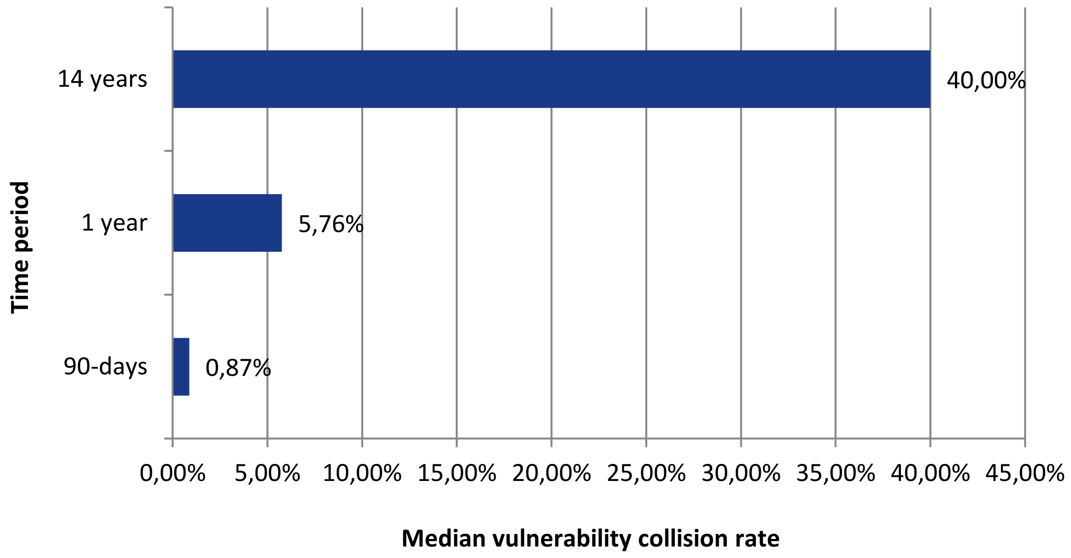

Collision rates are also featured in discussions of optimal disclosure times (i.e. how much time is allowed for a disclosure process before the vulnerability is publicly released). If vulnerability collision is zero, a disclosure process could, in theory, be as long as required to develop and roll out remediation measures to all affected parties. In contrast, if vulnerability collision is non-zero, there needs to be a discussion of how long a disclosure process can run before the vulnerability is disclosed (even if patch development or remediation work is outstanding). [32] The longer a CVD process goes on, the higher the risk of information leaks, independent rediscovery or the vulnerability being exploited by a malicious actor.

Vulnerability exploits can also be inexpensive and quick to develop once a vulnerability has been identified. The cost to develop an exploit is influenced by several factors: the time taken to identify the vulnerability; the time to develop the exploit; and the cost of purchasing specialist equipment to develop the exploit (e.g. particular code, infrastructure or testing equipment). The severity of the vulnerability may also influence the cost of the exploit development, but is more likely to affect the reward or price of the exploit (if it is sold). The uniqueness of the vulnerability may further influence the value of a vulnerability or exploit (e.g. if it is the only vulnerability identified in a particular product or if it meets the need of a particular customer or vendor).[33] Lastly, the location of where the vulnerability information is posted may also affect the exploit uptake (i.e. the more popular the forum where a vulnerability or exploit is published, the more it is used).[34]

The time required to develop a functioning exploit for an identified vulnerability is typically short: previous research from the RAND Corporation has shown a median time of 22 days to develop an exploit for zero-

[31] Ablon & Bogart (2017).
[32] Interview 11.
[33] Ablon & Bogart (2017).
[34] Rashid (2018).

day vulnerabilities.[35] In summary, this means that vulnerabilities are frequent, relatively inexpensively and quick to exploit, and subject to a comparatively long lifespan – where exploits remain active and useful for an extended period of time – which may have consequences for how vulnerability disclosure should be approached.

The following Chapter features a discussion on the economic aspects of the information security market and how this relates to vulnerability disclosure.

[35] Ablon & Bogart (2017).

3. Introduction to information security economics

A commonly held view is that information security ultimately comes down to technical measures. The reality, however, may be more complex.[36] There are several underlying factors that contribute to the persistent nature of vulnerabilities in hardware, software and services. While some are related to the nature of vulnerabilities themselves and their inherent characteristics, a number of others stem from the economic features of information security.

An analysis of the economics of vulnerability disclosure therefore requires an understanding of the underlying economic concepts of the information security environment. Many of these concepts stem from classical economics but have a particular relevance and application to information security. Section 3.1 of this Chapter discusses the following economic concepts and their relation to the information security market:

- Tragedy of the commons
- Network effects
- Externalities
- Asymmetric information and adverse selection
- Liability dumping
- Moral hazard.

Finally, Section 3.2 contains a brief discussion on how technological change affects the information security market and its associated economic attributes.

3.1 The information security market

Information security is often perceived to be subject to a **'tragedy of the commons'**.[37] A concept first introduced in 1833,[38] a tragedy of the commons refers to situations where individuals use a shared resource to the detriment of the collective good of all users of that resource. In this context, rational decision makers are expected to consider their own personal outcomes, or 'payoffs', without considering the payoffs of anyone else. Hence, assuming that individuals act only in their immediate self-interest, shared resources run the risk of being damaged or depleted to the long-term detriment of all parties.

In relation to information security, this concept can be seen in practice in the case of distributed denial-of-service (DDoS) attacks that utilise large numbers of compromised devices to generate enough traffic to render services unavailable. Whereas end users may be affected by the attack (e.g. not being able to access a website they would like to access), they are neither the primary target of the attack nor responsible for any of the costs of protecting against or restoring service after the attack (which typically fall on the service providers). In theory, these types of attacks could be rendered less effective if the end-user devices were of sufficient security. In an ideal world and functioning market, end users would secure their devices for the benefit of all other users and the Internet itself. However, without a personal incentive, such as cheaper access to the service or a regulatory requirement, individuals have very little desire to pay for something that will ultimately help others far more than it would help them individually. In relation to vulnerability disclosure, these dynamics also extend to developers and manufacturers who, in

[36] Anderson (2001).
[37] Bailey & Tierney (2002).
[38] Hardin (2009).

the absence of liability for vulnerabilities found in their products or services, may not be incentivised to engage in CVD or bug bounty programmes.

3.1.1 Network effects mean that products and services are subject to different security considerations in relation to their size and use

Network effects refer to the idea that as a greater number of people use a good or a service, the greater its value becomes to other users of that good or service.[39] Consider telephones: there is clearly no use in having a telephone if no one else has one, as one cannot use it for its primary purpose – communication. However, as more people start to buy telephones they become more and more valuable to both you and other telephone users, as there are more potential telephone owners to connect with.[40]

Networks effects like these can also be seen in the information security market. As a growing number of people use a certain software, product or service, many benefits are accrued as a result: further resources are devoted to its development and maintenance; more information becomes available about it; and more resources are dedicated to its security. Nevertheless, network effects can also be negative. A larger network will ultimately also become more complex, with more nodes and connections and a greater potential attack surface. A more prominent product or service may also attract greater attention from malicious actors, thereby offsetting some of the previous security gains.

As such, a portion of the value of a piece of software or service lies with how many users it has. However, two additional factors may also influence the value of software. First, following general microeconomic logic, hardware or software development typically has high fixed costs and decreasing marginal costs (e.g. the first output is costly to develop but subsequent outputs may be produced with decreasing marginal costs per unit produced). Second, technology markets are often characterised by high costs and other barriers for users when switching technologies or services, which leads to so-called ‘lock-in’ situations with dominant providers. The result is that these types of markets heavily benefit early entrants, where reduced time to market may bring significant commercial advantage and positive economic feedback loops – which may lead to firms prioritising speed to market over security, ultimately leading to less secure products or services.[41]

3.1.2 Externalities may have significant consequences for behaviour in the information security market

Another economic concept that underpins the information security market is that of **externalities.** In economic terms, an externality is a consequence of a market transaction (the buying and selling of a product or service) which impacts a third party (i.e. not the buyer or seller), without incorporating this effect into the market price.[42] A very simple example of an externality is driving a car. Owning a car has obvious **internal costs** to the driver, such as buying the car and paying for insurance or petrol, but there are also **external costs** that impact actors who are not involved in the purchase and sale of the car. These external effects include greenhouse gas emissions, lower air quality and increased congestion – all of which have a wider impact on society.[43]

Externalities are ubiquitous in the information security market. Technology is an integral part of today’s way of life and it is challenging to account for all of the potential wider implications of software or

[39] Anderson & Moore (2006).

[40] Anderson & Moore (2006).

[41] Anderson (2001).

[42] Van Eeten & Bauer (2008).

[43] Van Eeten & Bauer (2008).

hardware beyond its sellers and buyers. Software markets are typically competitive and relatively global, with many competing vendors in each market segment. Therefore, firms predominantly focus on their internal costs to remain competitive (e.g. costs of development, marketing, etc.), which may mean that they are incentivised to reduce costs for 'non-essential' features such as security.

Unlike many other goods, software and software components in hardware can also be fixed after release or time of sale. This means that vendors may be inclined to prioritise minimising internal costs and speed to market, rather than upfront software quality and security, with the intent to finalise functionality and security of the product at a later stage.[44] However, the deployment of insecure products or systems can have significant societal impact, including direct economic losses, data protection, privacy or reputational damages. These types of externalities also have direct effects on vulnerability disclosure (see Chapter 4).

3.1.3 Liability dumping may enable actors to shift responsibilities and costs to other actors

A very closely related concept to externalities is **liability dumping** – the idea of shifting the burden of responsibility onto other parties within a market.[45] Part of the reason that many vendors persist in developing and releasing insecure software is that, more often than not, they are not liable to bear the full cost of the consequences of their vulnerabilities. With several different actors involved in the market, there is no clear consensus as to who bears responsibility for what.

If a security incident occurs, it is unclear who should respond to it and who should be responsible for the costs of remediating the vulnerability. Should it be the government that takes responsibility, the software producers, the users of the product, or someone else? If actors in the information security market are not held liable for poor security – whether it is in relation to the security of the software, the security of an individual's device or a corporate network – they are less likely to invest in adequate security measures, particularly if the costs are borne elsewhere.

Liability dumping or transfer is often related to cost. In relation to vulnerabilities, two types of costs can be incurred if a vulnerability is identified: the **potential damage cost** that would be incurred if the vulnerability is exploited; and the **patching cost** associated with identifying, testing and rolling out an appropriate remediation measure. While vendors are expected to release appropriate remediation measures, which will incur some costs internal to the vendor, they are rarely held accountable for damage costs that other actors may incur in the event of the exploitation of the vulnerability.[46]

3.1.4 A lack of or variability of knowledge results in sub-optimal choices in the information security market

Part of the difficulty in incentivising security at different stages of the information security market is due to the presence of **asymmetric information** and **adverse selection**. Asymmetric information exists where either the buyer or the seller in a given market knows more about the good or service for sale than the other. This may lead the less informed actor to pursue sub-optimal choices in their transaction, which ultimately affects the overall quality of the market (due to adverse selection).[47]

In classical economics, this has been illustrated by the used car market. For example, consider a used car market where half of the cars are of high quality and half are of low quality, and only the sellers know which are which. Since buyers are unwilling to pay high prices for fear of getting a low-quality car in return,

[44] Arora & Telang. (2005).
[45] Anderson, Böhme, Clayton & Moore (2008).
[46] Cavusoglu et. al. (2006).
[47] Anderson (2001).

the market price for all cars decreases. However, sellers with high-quality cars may not be willing to sell the cars at a lower price, prompting them to exit the market and leaving it predominantly populated by low-quality cars.[48]

This type of adverse selection, driven by information asymmetries, is also observed in the information security market.[49] General users may make decisions to use or buy software, hardware or services influenced by a number of considerations, including:

- Price
- Ease of use and convenience
- Security
- Value of the service offered.[50]

Trade-offs between price, convenience and security often arise. Users may value security considerations more highly than price or convenience in relation to services they perceive as sensitive, such as banking services, while they value convenience more highly than security for day-to-day services, such as general smartphone apps.

For marketing purposes, vendors often make claims as to the level of quality of the security of their product or service, but customers cannot typically evaluate the accuracy of those claims, either because of a lack of information or a lack of the technical knowledge required to make an informed assessment.[51] As a result, customers may be unwilling to pay for a product advertised with additional security features but at a higher price relative to other similar products. One solution for this particular problem could be product certification or labelling as a way to offset the disadvantage.

Users are generally content to overlook security practices as long as their incentives – financial or otherwise – exceed their inconvenience. The relative lack of importance placed on security by consumers may therefore lead to negative incentives for vendors and service providers, who may instead prioritise convenience or price over security. The trend may contribute to a market where vendors are not incentivised to invest in increased security for their products, crowding out more secure software and leaving a market of lower quality, more vulnerable products.[52]

3.1.5 Moral hazard means that decision makers in the information security market take decisions subject to high risks

Asymmetric information is also a pre-requisite for a **moral hazard** to exist – a situation where one party 'makes the decision about how much risk to take, while someone else bears the cost if things go badly'.[53] A classic example of a moral hazard is house insurance, where an insured individual may behave more recklessly than an uninsured person because, if something does happen, the insurance company is the party that will bear the costs incurred (rather than the individual).

In the information security market, actors may be incentivised to accept more risk than is socially optimal due to the perception that others may bear the consequences should the risk be realised. Actors may also

[48] Akerlof (1970).
[49] Anderson (2001).
[50] HackerOne (2016).
[51] Interview 12.
[52] Anderson et al. (2009).
[53] Krugman (2009).

accept additional risk simply because they cannot understand or determine what the socially optimal security decisions would be. Previous research from the University of Maryland has shown that average users face significant challenges in determining optimal security strategies in situations where interdependencies between security choices of other actors are present, which may lead to security strategies that fail.[54] For example, a corporate user may open a link in a suspicious email out of curiosity, thinking that corporate security policies will protect them, not fully appreciating the risk that can be realised to other users within the organisation.

3.1.6 The economic features of the information security market affect the security of products and services, but also influence actors within the vulnerability disclosure landscape

This chapter has thus far illustrated that information security is a complex endeavour. The economic features of the information security market frequently result in perverse incentives and less than optimal economic behaviour, which contributes to the insecurity of many products and services, as well as the persistent nature of vulnerabilities.

The economic concepts discussed in the preceding sections are also directly relevant to vulnerability disclosure. As vulnerability disclosure takes place in the information security market, economic features present in this market also feature in the economic considerations of vulnerability disclosure. These economic concepts and features may also have specific consequences for particular parts of vulnerability disclosure processes. For example, organisations may be less incentivised to implement due to externalities that mean that they are not held liable for vulnerabilities found within their products (see Section 4.4.2), or individual security researchers may be unjustly held accountable for vulnerabilities due to liability dumping through unfair legal contracts between vendors and researchers (see Section 4.3.2). Chapter 4 discusses the economic considerations of vulnerability disclosure in further detail and highlights when and how the economic features of the information security market translates into vulnerability disclosure.

3.2 The changing nature of the information security market

The information security market also evolves alongside technological developments; the information security market of yesteryear is not the market of today. The combination of a number of factors continuously shape the market conditions and interactions, including decreasing costs of computing power, increasing levels of computing complexity, and increasing levels of connectivity and connected devices.

Today, software and hardware is embedded almost everywhere and used by almost everyone at all times. Not only is society increasingly connected through the Internet and other networks, but it is also populated by individuals and organisations with an increasing number of connected digital devices. The consequences for vulnerability disclosure are two-fold. The increased use of technology results in a greater possible vulnerability surface – particularly considering when technology is deployed or embedded in new environments. Most individuals may have a basic understanding of possible vulnerabilities in a computer or smartphone, whereas many individuals, organisations and policy makers only have a nascent understanding of what security means for Internet-connected fridges, connected medical devices, or other connected devices that have significant computing capabilities.

An increased development and use of connected applications, technologies and services by organisations that have not traditionally had a significant ICT presence or development function may further jeopardise

[54] Gordon & Loeb (2002); Acquisti et al. (2005).

overall ecosystem security.[55] The automotive sector is an example of an industry that traditionally did not develop or integrate computers or connectivity within their cars, but which is now increasingly developing cars with an array of integrated software and hardware components, many of which are third-party solutions. Therefore, car manufacturers have to familiarise themselves with the current state of the art in secure development and information security, as well as develop an understanding of how to design, implement and operate vulnerability disclosure policies. This may require a technical and organisational maturity that many new entrants might currently lack.

The increasing levels of digitisation and connectivity are in part due to the decreasing costs of computing power and the increasing levels of computing complexity.[56] Complex systems are comparatively more challenging to secure than simpler counterparts, and complex systems are becoming increasingly common.

While the future of Moore's law may be uncertain, Central Processing Unit (CPU) complexity is still increasing.[57] Historically, it was typically more expensive to build a complex machine compared to building a simple one. In modern computing, this may not be true and it is often more cost-effective to use a complex CPU and simulate simplicity, rather than to develop a simple device from scratch – as complex general-purpose CPUs are readily available and inexpensive. In practice, software is added on top of hardware complexity to simulate simplicity (i.e. to enable the device to perform the required functions), which also means that one can add additional features later on in the process (whether intended or not) at near-zero short-term marginal cost.[58] The use of low-cost, complex CPUs in the production of simple devices – which also increases the complexity in securing the device – can be referred to as an 'anomaly of cheap complexity'.

The anomaly of cheap complexity is further exacerbated by shifting development and manufacturing approaches, as vendors continue to prioritise time to market, as shown in Figure 3.1.

Figure 3.1 Shifting development lifecycles

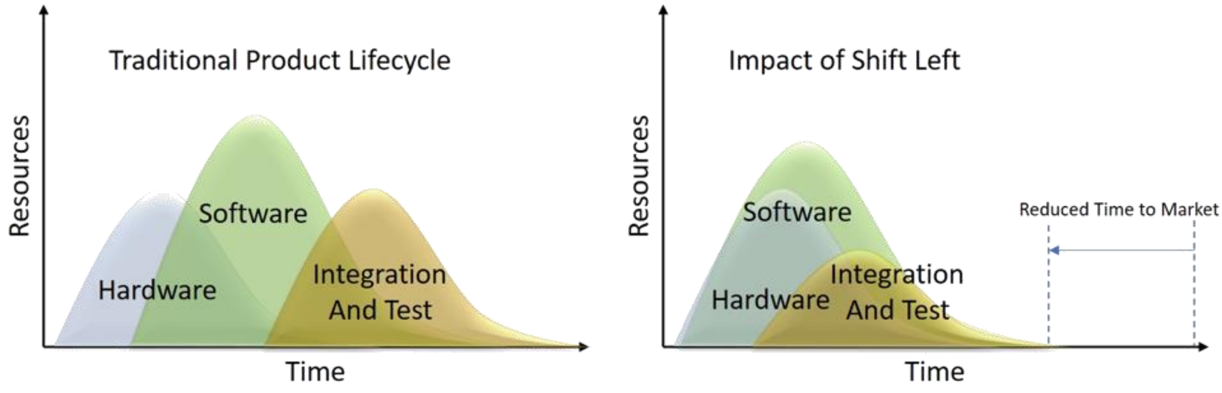

Source: Dullien (2018).

[55] Interview 5.
[56] Dullien (2018).
[57] Moore's law refers to the prediction that the number of transistors in a dense integrated circuit is expected to double around every two years.
[58] In his conference proceedings, Dullien presents the example of an Apple Lightning to DVI adapter that has an integrated CPU with the equivalent computing power to a PlayStation confined in what is essentially just a cable.

The anomaly of cheap complexity, increased connectivity and shifting development and manufacturing priorities can result in a number of inter-related network effects and security challenges. The use of general purpose CPUs introduces a level of complexity that may not be required for the device to function as expected, but ultimately makes security more difficult to achieve. The inherent complexity of many of these devices may enable an attacker to perform a function not originally intended (e.g. using Internet-connected security cameras for a Distributed Denial Service (DDoS) attack as in the case of the Mirai botnet).[59] It may also make it more difficult to inspect and asses the level of security that the device currently exhibits by making security testing prohibitively complex, impractical or expensive. The global distribution of products and services, together with an increased focus on reducing the time to market, may also result in unintended and potentially harmful consequences. If the speed of development takes precedence over security, insecure codes or new codes with vulnerabilities may be embedded into products and services. This may not be a significant security problem in isolation, but if the product is an inexpensive Internet of Things (IoT) device that is sold in millions of units, an exploitable vulnerability in that product may have severe security implications.

These security challenges are particularly troublesome in light of current and predicted growth in IoT devices and trends in cyber-physical devices.[60] This is partly due to the likelihood that vulnerabilities in cyber-physical systems may have real-world effects,[61] and partly due to the fact that many of these types of systems are expected to be in operation for an extended period of time with limited intervention (e.g. infrequent security assessments or patches). The rapid growth in IoT and cyber-physical devices has resulted in a wide range of complex policy challenges in terms of how current legislation, liability, insurance, consumer safety and other policy regimes will have to be revised to be able to meet these emerging security challenges.[62]

The changing nature of the information security market also has consequences for the economics of vulnerability disclosure, how vulnerability disclosure is carried out, and how actors behave in the vulnerability disclosure process – some of which are only now beginning to be discussed and understood. As with traditional development and manufacturing approaches, the current thinking of vulnerability disclosure grew out of a traditional, phased development approach where software and hardware were finalised and released in distinct releases (e.g. ‘shrink-wrapped software’).[63] New approaches to development and new types of embedded, IoT or cyber-physical devices (also referred as Industry 4.0) of increasing complexity may warrant examination of current practice to understand if it is still fit-for-purpose, efficient and effective.

[59] The Mirai botnet took advantage of poorly secured IoT devices, which were compromised through Telnet connections enabled by a list of common login credentials, to launch DDoS attacks. See https://www.cloudflare.com/learning/ddos/glossary/mirai-botnet/ (As of 31 October 2018).

[60] Cyber-physical systems refer to integrations of computation, networking and physical processes (i.e. a connection between the physical and the virtual worlds). This typically takes the form of embedded devices that monitor and control physical processes in devices such as connected and intelligent traffic monitors or wearable, connected healthcare devices.

[61] In the words of computer scientist Ross Anderson of Cambridge University: ‘Phones and laptops don’t kill many people directly: cars and medical devices do’ Anderson (2018).

[62] Many of which are discussed at length in Leverett et al. (2017).

[63] Interview 12.

Recent years have already seen an increase in the number of high-profile vulnerabilities requiring extensive multi-vendor coordination, such as Heartbleed, Meltdown/Spectre and others.[64] It is feasible that these types of vulnerabilities that have significant downstream effects or security implications will be increasingly common, and the vulnerability disclosure community may need to examine if current processes adequately address the type of harm that such vulnerabilities could result in and how coordination of these types of vulnerabilities should be performed.[65]

The economic considerations and incentives that influence the behaviour of the different actors within the vulnerability disclosure process are further discussed in the following chapter.

[64] Interviews 5, 9, 12. For a discussion of the Heartbleed vulnerability see ENISA (2016), and for a discussion on Spectre-Meltdown see Chapter 5.
[65] Interviews 5, 9, 12.

4. Incentives and behaviour in vulnerability disclosure

4.1 Introduction

The development of secure hardware, software and other services is subject to numerous challenges. The demand for feature-rich solutions in competitive and fast-moving markets further exacerbates the complexity of developing such systems, particularly when considering the need to manage legacy code bases or hardware and interdependencies with other systems.[66] While adoption of new development practices, such as DevOps[67], and advances in automated security testing technologies may improve overall system security, it is unlikely that vulnerabilities will entirely disappear. Instead, organisations and vendors are increasingly utilising external security researchers to crowdsource efforts to identify and remedy vulnerabilities, so as to mitigate the overall security (and commercial) risk.[68]

The vulnerability disclosure process is dictated by the actions of multiple economic agents, many of whom may have competing or conflicting interests. As a result, sub-optimal or even undesirable outcomes of the kinds outlined in section 3 are pervasive. Through different regulations or other incentive mechanisms, it is nevertheless possible to change or influence these outcomes. In order to better understand the economics of vulnerability disclosure, it is helpful to map some of the underlying economic considerations and incentives that may influence behaviours across the disclosure lifecycle.

This chapter begins with a brief discussion on the nature of incentives (Section 4.2). It further discusses these economic considerations and incentives across four different levels:

- **Individual level**, comprising finders (Section 4.3)
- **Organisational level**, comprising vendors, coordinators and governments (Section 4.4)
- **Structural level**, relating to regulatory or legal incentives (Section 4.5)
- **Normative level**, relating to social norms, good practice and voluntary standards (Section 4.6).

4.2 The nature of incentives

Economic decisions in relation to vulnerability disclosure are shaped by incentives perceived by different actors at the different decision points in the vulnerability disclosure lifecycle. Incentives in this regard can be understood as the factors that influence these decisions. Incentives may stem from the economic features of information security, or other legal, regulatory or market conditions, as well as originating from socially-expected behaviour and norms.[69]

It is important to note that economic incentives are not only connected to financial considerations or the related costs of a specific vulnerability disclosure. In contrast, incentives are typically divided into financial and non-financial incentives. As seen in the previous chapter, financial considerations are pervasive in the information security market and are fundamental drivers of vendor behaviour in relation to securing market shares and maximising profitability. In relation to vulnerability disclosure, financial incentives can,

[66] Anderson (2002); Hahn & Layne-Farrar (2006).

[67] DevOps is a software development practice that aims to integrate development and operations teams in order to work in more agile and shorter production cycles. DevOps may also enable better integration of security teams in the development process and therefore contribute to proactive, rather than reactive, security. For more information on DevOps practices, see Rahman & Williams (2016).

[68] Finifter et al. (2013); Zhao et al. (2015).

[69] Van Eeten & Bauer (2008).

for example, take the form of monetary rewards for finders that identify valid vulnerabilities as part of bug bounty programmes. Financial incentives may also influence the vulnerability disclosure process if actors motivated by financial incentives choose not to disclose an identified vulnerability and instead sell it through a vulnerability broker or marketplace.

In contrast, non-financial incentives are considerations or characteristics that influence behaviour that is not directly connected to money. Non-financial incentives may include considerations related to socially-constructed norms or values, where actors behave in a particular way based on ethical ideals or where behaviour is self-regulated in line with perceptions of expected behaviour. Financial incentives therefore typically reward achievement or results (negative or positive), whereas non-financial incentives usually work through self-regulation or peer pressure.[70]

Incentives may also interact with each other by complementing, trading off or contradicting one another. In order to better understand the economics of vulnerability disclosure, it is therefore necessary to explore how the behaviour of different actors is incentivised and how the different incentives interact with each other.

4.3 Finder behaviour

Finders are necessary for any vulnerability disclosure process. Finders may be located within a vendor, be external to the organisation but hired by a vendor, or be external and independent to the vendor. The latter two categories comprise vulnerability research teams and organisations[71] and independent security researchers.[72]

4.3.1 Internal incentives act as the primary driver for finder behaviour in the vulnerability disclosure process

As illustrated in Figure 4.1, previous research and interviews conducted as part of this study reveal that finders generally engage in vulnerability research and disclosure due to four overall categories of motivations and incentives.[73]

[70] Van Eeten & Bauer (2008).

[71] Such as academic research groups or corporate vulnerability research teams, for example Google Project Zero.

[72] These researches may perform vulnerability research as part of or on top of their normal day job.

[73] I Am the Cavalry (n.d.), HackerOne (2017), Interviews 1, 2, 8, 10.

Figure 4.1 Incentives and motivations of finders

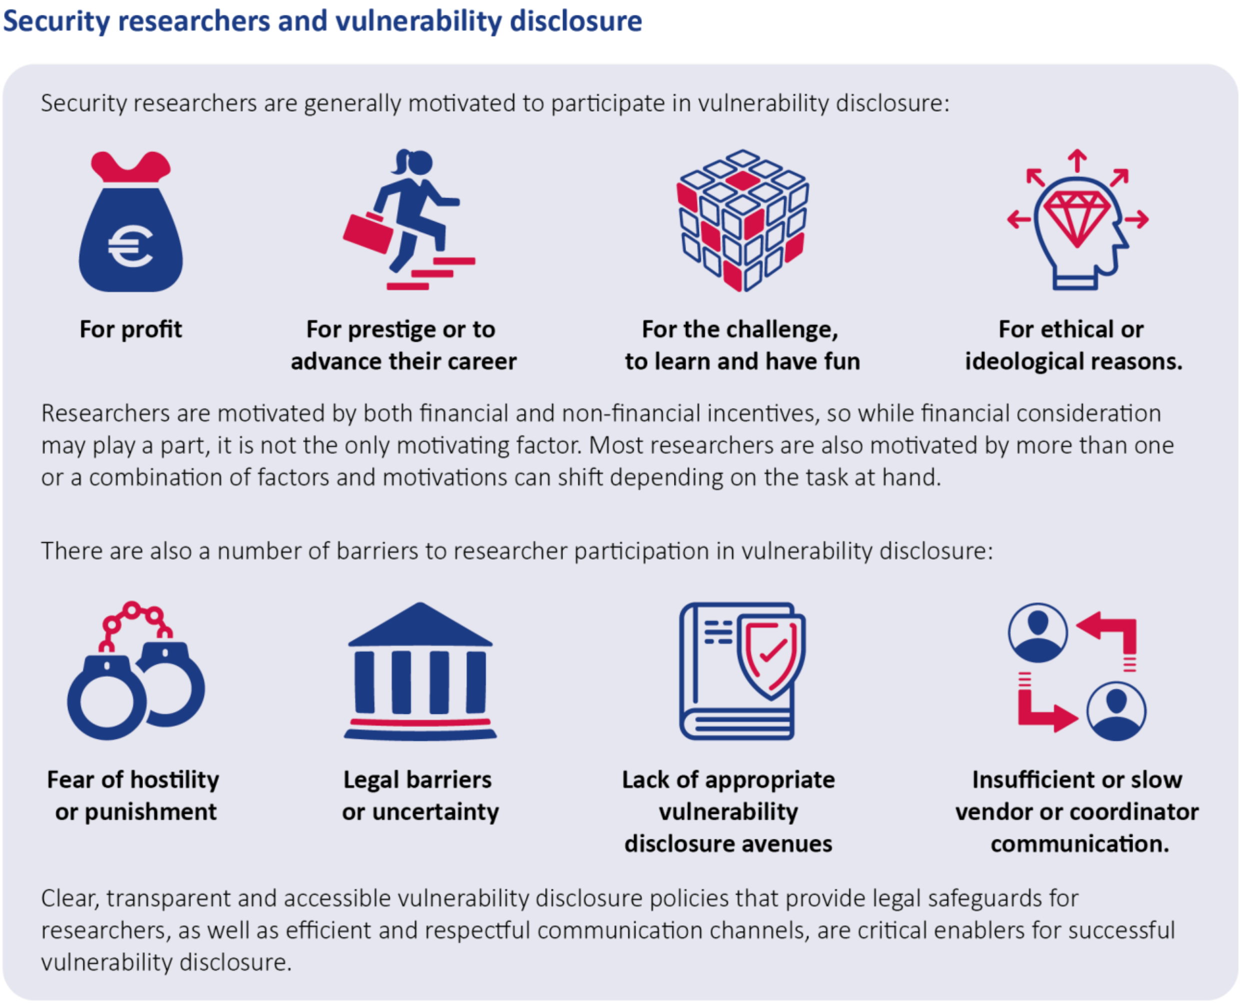

Security researchers and vulnerability disclosure

Security researchers are generally motivated to participate in vulnerability disclosure:

For profit

For prestige or to advance their career

For the challenge, to learn and have fun

For ethical or ideological reasons.

Researchers are motivated by both financial and non-financial incentives, so while financial consideration may play a part, it is not the only motivating factor. Most researchers are also motivated by more than one or a combination of factors and motivations can shift depending on the task at hand.

There are also a number of barriers to researcher participation in vulnerability disclosure:

Fear of hostility or punishment

Legal barriers or uncertainty

Lack of appropriate vulnerability disclosure avenues

Insufficient or slow vendor or coordinator communication.

Clear, transparent and accessible vulnerability disclosure policies that provide legal safeguards for researchers, as well as efficient and respectful communication channels, are critical enablers for successful vulnerability disclosure.

Source: ENISA study on the Economics of Vulnerability Disclosure

Financial incentives play a particularly prominent role in bug bounty programmes, where they are regularly used to incentivise participation from the finder community. While not all data on bug bounty programmes is public, and despite the fact that a large number of bug bounty programmes are private, public reports from the two largest bug bounty platforms, Bugcrowd and HackerOne, can provide a glimpse into the state of bug bounty programmes globally. According to Bugcrowd, the average bounty paid out per vulnerability on their platform in 2018 was US$781, which corresponds to a 73 per cent increase, compared to 2017. The average bounty for the most severe category of vulnerability was US$1,200, compared with US$926 the year before.[74] In total, Bugcrowd paid out more than US$6 m in bounties from 2017–2018.

Both the total volume of vulnerabilities reported submitted and the total amount of bounties paid out has, in recent years, increased on both platforms. Bugcrowd experienced an overall increase of 40 per cent in bug bounty programmes launched in 2018, compared to 2017, while HackerOne has seen a 10-fold increase in registered finders in just two years.[75] This growth may also enable further positive network

[74] Bugcrowd (2018).
[75] Bugcrowd (2018); HackerOne (2018).

effects, where bug bounty platforms can both leverage economies of scale and help raise the profile of vulnerability disclosure across the information security market. Previous research from Pennsylvania State University has also shown that financial incentives have a significant positive correlation with the number of reported vulnerabilities.[76] The presence of financial incentives has also been found to have a correlation with the quality of vulnerabilities identified, as bug bounty programmes without monetary rewards may receive more vulnerability reports but of lesser quality.[77]

The increased use of bug bounties has also been perceived to contribute to the 'professionalisation' of the finder community and to the development of a new generation of information security professionals.[78] HackerOne highlights that many of their most active finders use bug bounty programmes for part of their monthly income. Some 25 per cent of researchers surveyed by HackerOne rely on bug bounties for at least 50 per cent of their annual income, and over 13 per cent state that bounties represent 90 –100 per cent of their annual income. Around 12 per cent of researchers on HackerOne make US$20,000 or more annually from bug bounties, with over 3 per cent making more than US$100,000 per year.[79] These annual income levels may seem low, but many of the researchers active on Bugcrowd and HackerOne live and operate in countries outside the European Union where living costs are lower, such as India, Pakistan and Russia.[80]

In wider information security economics, security investment is often regarded as having positive but diminishing returns. However, this does not seem to be the case in relation to vulnerability disclosure. If diminishing returns were present, disclosure and fix rates should decrease and the time between vulnerabilities would increase as vulnerabilities would become fewer and more difficult to remedy. In contrast, it seems that vulnerability disclosure or fix rates are not declining and that the supply of vulnerabilities that can be patched seems to be steady.[81] This can also be seen in the fact that many exploited vulnerabilities are commonly known or have been known for a long time, indicating recurrent patterns in human error or inadequate development practices.[82]

However, diminishing returns for individuals may influence finder behaviour in bug bounty programmes – particularly if the finder is primarily motivated by financial incentives (and therefore tries to maximise financial rewards while minimising their effort).[83] The longer an individual bug bounty programme has existed the more likely it is that easily identifiable vulnerabilities have been submitted, and remaining vulnerabilities are more complex or difficult to find. Activity levels and the rate of enrolment of new finders tend to decline over time and finders may instead turn their attention to newly established bug bounty programmes with more easily identified vulnerabilities (and thus more immediate financial and non-financial returns).[84] Nevertheless, interviewees noted that this is mostly in relation to profit-incentivised security researchers who tend to rely on automated tools and submit low quality vulnerability reports.[85]

[76] Zhao et al. (2015).
[77] Ruohonen & Allodi (2018).
[78] Interviews 1, 2, 9.
[79] HackerOne (2018).
[80] HackerOne (2018).
[81] Neuhaus & Plattner (2013).
[82] Verizon (2018).
[83] Zhao et al. (2017).
[84] Maillart et al. (2016).
[85] Interviews 4, 13.

Nevertheless, the predominant motivation for most finders to participate in vulnerability disclosure is for the technical challenge it offers, as well as to prove their abilities compared to other security researchers.[86] The information security community is typically highly competitive, where finders pride themselves on their standing vis-à-vis their peers.[87] Popular bug bounty programmes make extensive use of point systems and leaderboards where researchers can compete for standings and other achievements.[88]

In contrast to a 'tragedy of the commons', where economic actors simply consider their own payoffs above others, many finders also participate in the vulnerability disclosure process due to personal ethics or a sense of duty to make a positive difference in the researcher community and to help improve overall information security.[89] In some cases, this sense of duty may also carry more importance than financial incentives, as seen with finders donating their bug bounties to charitable causes.[90] This type of behaviour can also be seen to extend beyond the traditional security researcher community. One interviewee highlighted that many of the vulnerability reports received by their organisation were submitted by general users of the service rather than security researchers, particularly in relation to business logic vulnerabilities. These finders were also found to rarely ask for or accept bug bounties for their reports.[91]

The strong ethical foundation within the security researcher community often results in strong self-regulation of behaviour throughout the vulnerability disclosure process. This type of self-regulation may manifest itself in several ways. Sometimes this is outwardly facing, with security researchers expressing concern or appreciation for the way vulnerability disclosure is carried out by vendors or coordinators.[92] Historically, many security researchers argued that only full disclosure was the ethically correct choice and the only way that security could be improved.[93] Today, by contrast, many security researchers argue to follow a CVD process in order to allow vendors to develop and roll out appropriate remediation measures, whether this involves rewards or not.[94]

The available data on finder motivations should be taken into account when designing a CVD or bug bounty programme. Particularly, it should be clear that organisations can benefit from the wider security research community by operating only a CVD programme without financial compensation to researchers.

4.3.2 Finders also respond to a number of external economic considerations and incentives

The discussion has thus far focused on incentives that are internal to the finder, stemming from choices and motivations that characterise the finder. However, there may be additional influencing factors that are beyond the direct control of the finder but that nevertheless influence their behaviours and choices in the vulnerability disclosure process. These factors include:

- Fear of hostility or punishment
- Legal barriers or uncertainty
- Existence of appropriate vulnerability disclosure avenues

[86] Interviews 1, 2, 8, 10.
[87] HackerOne (2017).
[88] See, for example, https://hackerone.com/leaderboard/invites and https://www.intigriti.com/public/leaderboard (As of 31 October 2018).
[89] Interviews 2, 8, Householder et al. (2017); Lewis (2017).
[90] HackerOne (2017).
[91] Interview 6.
[92] Lewis (2017).
[93] Arora & Telang (2005).
[94] Zero-day Initiative (n.d.).

- Presence, efficiency and quality of vulnerability disclosure processes
- Structure and quality of the vendor or coordinator communication mechanisms.

Vendors may not possess the adequate technical skills, interest or resources to receive and remedy identified vulnerabilities that are received from finders. Finders or 'hackers' have traditionally been subject to discrimination and suspicion as to whether their motivations are pure or whether they are acting maliciously. As such, many finders have reported vulnerabilities only to realise they are not listened to, or else are met with hostility from vendors, including threats of prosecution.[95] A vulnerability disclosure landscape that places significant **fear of punishment** on finders may therefore have adverse effects on the number and quality of vulnerabilities identified, disclosed and ultimately mitigated (i.e. act as a negative incentive and deterrent).

Moreover, security researchers engaging in vulnerability disclosure often move in **legal grey areas**, particularly when there are no established processes for vulnerability identification or disclosure.[96] Unauthorised access to or control of software, hardware or services is often illegal – even in the presence of a valid, identified security vulnerability, sometimes even when there is a CVD policy in place.[97] Legal implications for security researchers may extend to both civil and criminal law, as well as contract law, licensing, patent law and other types of legislation.[98] If legal safeguards are not provided for finders, vendors may shift liability for discovered vulnerabilities to security researchers in order to avoid being held accountable for any costs incurred by the vulnerability, even if the vulnerability was reported in good faith (i.e. liability dumping).

In the absence of clear regulatory and legislative regimes for vulnerability disclosure, organisations are left to create alternative regimes to enable security research through market mechanisms and standard form contracts.[99] A standard form contract is a contract between two parties in which the terms and conditions of the contract are defined by one of the parties, and the other party has little or no ability to negotiate more favourable terms. In CVD, and particularly in bug bounty programmes, legal boundaries and contract terms are typically dictated by the organisations or vulnerability platforms, leaving little room for individual security researches to negotiate or change the terms of the disclosure – in essence forcing the security researchers into a 'take-it-or-leave-it' position.[100] Many security researchers also lack the legal expertise to accurately evaluate proposed contract terms,

This type of market regulation – through unilaterally drafted boilerplate language that functions as private law – highlights the agency problem faced by many security researchers. Most security researchers are not lawyers or educated on legal matters, making it difficult for them to evaluate the possible personal legal implication of participating in a particular vulnerability disclosure process, which can lead to suboptimal choices (i.e. adverse selection) driven by information asymmetries between finders and vendors. As contract terms are typically drafted to foremost protect the organisation and/or vulnerability platforms, the legal risks for disclosing vulnerabilities may be shifted onto the individual researcher (i.e. liability dumping), particularly in the presence of legal grey areas related to anti-hacking laws within the relevant jurisdiction.[101] This presents significant challenges within the vulnerability disclosure landscape, as one

[95] Interviews 1, 8, 9, Lewis (2017).
[96] Interviews 1, 8, 13.
[97] Interview 13.
[98] Peeters (2017).
[99] Interview 13.
[100] Interview 13.
[101] Elazari Bar On (2018).

study found that as much as 60 per cent of security researchers cited the threat of legal repercussions as a reason they might not work with a vendor in the disclosure of a vulnerability.[102]

As security researchers have little to no choice in negotiating contract terms, and there is still reasonable participation from the security community in vulnerability disclosure (i.e. organisations continuously receive reports), there is currently little incentive for organisations to change their legal approach. If the security researcher community started to attach further reputational or economic value to contract terms, organisations may have to adjust their approach to attract satisfactory levels of participation in their vulnerability disclosure programmes. There are currently several initiatives, mostly in the United States, that aim to improve practice around legal safeguards for security researchers conducting security research in good faith.[103] These initiatives collaboratively aim to improve safe harbour[104] for researchers and programme owners and readability of legal terms. However, the diversity of the European legal landscape can make it challenging to realise this type of safe harbour, as Member States have different legal systems and different approaches to hacking-related legislation. [105] However, some European vulnerability disclosure actors have begun to enact similar clauses within their programmes.[106]

The presence of **appropriate vulnerability disclosure avenues** is thus imperative to incentivise positive finder behaviour. As discussed, vulnerability disclosure processes may be structured in several different ways, both internal and external to the organisation in question. Larger organisations with more resources available for security work may choose to develop internal vulnerability identification and disclosure mechanisms. Smaller organisations may be more likely to choose to enlist the help of external experts – either through outsourcing, contracting or public-facing vulnerability or bug-bounty programmes. As such, having established avenues for vulnerability disclosure may act as a positive incentive for finders – particularly if there are multiple avenues for engagement.

Beyond the existence of vulnerability disclosure avenues, **the institutionalisation and quality of vulnerability disclosure processes** also affect finder perceptions and willingness to engage with a particular vendor or coordinator. Vulnerability disclosure processes must be publicly available, easy to understand and provide clear guidelines as to the scope of the vulnerability disclosure process, its requirements and its expectations.[107] A clear and institutionalised policy may help finders alleviate concerns or fear of punishment when approaching a vendor for vulnerability disclosure.

Related to the vulnerability disclosure process, the **structure and quality of communication** is another potential factor that can positively influence finder behaviour.[108] Clear, secure and useful communication between the actors is paramount to a successful vulnerability disclosure process.[109] While vulnerability

[102] NTIA (2016).
[103] Including Dropbox's call to protect security researchers (https://blogs.dropbox.com/tech/2018/03/protecting-security-researchers/), GitHub's Open Source Vulnerability Disclosure Framework (https://github.com/bugcrowd/disclosure-policy), Amit Elazari's #legalbugbounty (https://github.com/EdOverflow/legal-bug-bounty), and the recently launched disclose.io project (https://disclose.io/) (As of 31 October 2018).
[104] In vulnerability disclosure, safe harbour refers to the practice of providing legal safeguards for security researchers participating in CVD or bug bounty programmes as an assurance that legal action will not be taken undertaken under anti-hacking laws for security research conducted in good faith.
[105] Peteers (2017).
[106] Interview 7. See for example https://www.intigriti.com/public/project/telenet/base (As of 31 October 2018).
[107] Interview 1, FIRST (2017).
[108] Interviews 1, 9, 10.
[109] NTIA (2016).

disclosure communication is typically initiated by finders, an oftentimes long and complex communication process follows in order to validate and remediate the vulnerability.[110] Finders expect this process to be marked by regular communication with the vendor or coordinator, and a failure to meet these expectations is often cited by finders as a leading reason for abandoning a responsible disclosure process.[111]

Effective communication in the vulnerability disclosure process can be facilitated by a number of measures, including secure channels for reporting, provision for anonymous reporting, provision of efficient communication with useful information, and prevention of premature disclosure of information.[112] Communication is key to maintaining trust in the vulnerability disclosure process, which is an important consideration made by finders when choosing whether or not to engage a coordinator or vendor.[113] Coordinators or vendors that are perceived to have a poor vulnerability disclosure track record or that have a history of unresponsive communication may be avoided by finders. Finder experiences with other actors in the vulnerability disclosure landscape are often shared with peers and the community, particularly when negative, and may in turn influence the decisions of other finders in the future.[114] As such, trust between the actors within vulnerability disclosure and a mutual recognition of the importance of protecting the shared ecosystem are key in order to avoid an information security ‘tragedy of the commons’.

4.4 Organisational considerations and incentives

Organisational considerations and incentives cover both private and public sector organisations, such as vulnerability coordinators or vendors, and government agencies acting on behalf of national interests.

4.4.1 Private and public sector organisations respond to regulatory and user expectations while operating with limited resources

As seen in Chapter 3, the economic features of the information security market influences the behaviour of vendors in the development and manufacture of software and hardware, sometimes resulting in negative outcomes that contribute to the persistent nature of vulnerabilities. The persistence of vulnerabilities pose challenges to information security and vulnerability disclosure efforts, particularly when faced with competing economic incentives such as speed to market versus security. Organisational decisions in vulnerability disclosure are therefore taken in the context of both the economic features of information security and the economic incentives, and considerations of vulnerability disclosure.

In vulnerability disclosure, private and public organisations can take on different roles, including as vendors or vulnerability coordinators. Vulnerability coordinators are guided by their mandate and their behaviour is typically aimed at maximising the social benefits of CVD. In contrast, the participation of vendors in the vulnerability disclosure process is often multifaceted and influenced by a number of economic considerations and incentives, as shown in

Figure 4.2.

[110] Lewis (2017).

[111] Interviews 1, 9, 10, NTIA (2016).

[112] Householder et al. (2017).

[113] Interviews 1, 5, 10.

[114] Householder et al. (2017).

Figure 4.2 Incentives and motivations of organisations

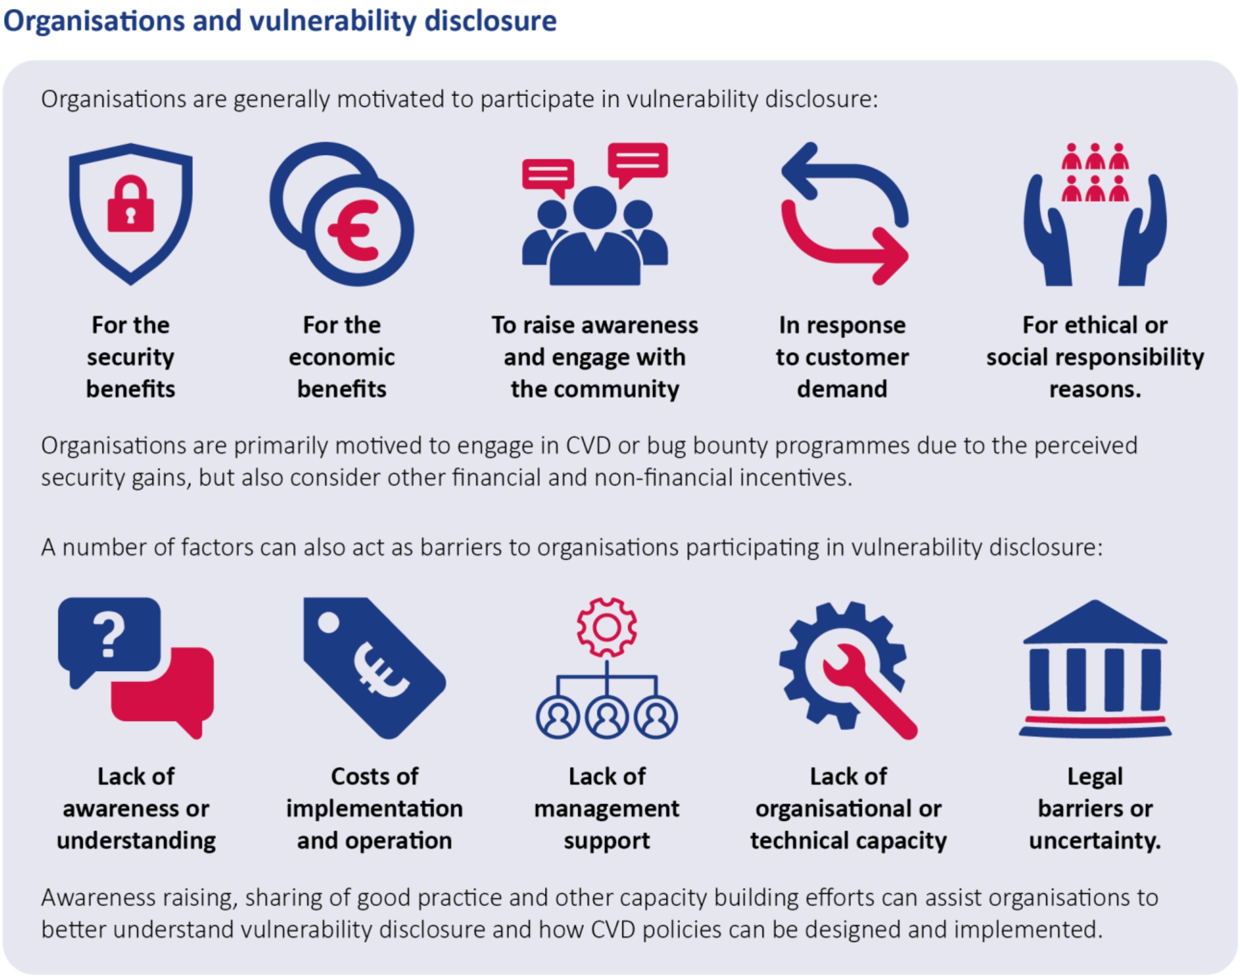

Organisations and vulnerability disclosure

Organisations are generally motivated to participate in vulnerability disclosure:

- For the security benefits
- For the economic benefits
- To raise awareness and engage with the community
- In response to customer demand
- For ethical or social responsibility reasons.

Organisations are primarily motived to engage in CVD or bug bounty programmes due to the perceived security gains, but also consider other financial and non-financial incentives.

A number of factors can also act as barriers to organisations participating in vulnerability disclosure:

- Lack of awareness or understanding
- Costs of implementation and operation
- Lack of management support
- Lack of organisational or technical capacity
- Legal barriers or uncertainty.

Awareness raising, sharing of good practice and other capacity building efforts can assist organisations to better understand vulnerability disclosure and how CVD policies can be designed and implemented.

Source: ENISA study on the economics of vulnerability disclosure

Organisations may be incentivised to engage in vulnerability disclosure for **economic benefits**. Vendors may perceive direct economic benefits from engaging in vulnerability disclosure by, for example, reducing development, marketing or security assurance costs. [115] CVD programmes enable organisations to realise efficiency gains from the ability to yield the effort and knowledge of a large number of security researchers for a relatively low effort and cost.[116] Bug bounty programmes may also be an effective option for both larger and smaller organisations, with research showing that such programmes can be from 2–100 times more cost-effective than hiring external security research to identify vulnerabilities.[117] However, CVD and

[115] NTIA (2016).

[116] Zhao et al. (2017).

[117] Zhao et al. (2015) ; Finifter et al. (2013).

bug bounty programmes may not be able to perform to the same scope or depth compared to penetration tests or other forms of security testing.[118]

Organisations are, however, perceived to primarily engage in CVD or operate bug bounty programme due to the anticipated **security benefits** that these types of programmes can bring – particularly in an environment where the goal is to identify vulnerabilities before a malicious actor does. In a competitive information security market where suitably qualified and experienced personnel are scarce and in which the attacker has a competitive advantage, crowdsourcing security through the use of CVD or a bug bounty programmes can help organisations shift the security balance in favour of defence.[119] The attractiveness in the possibility to mobilise a large number of security researchers again shows the potential of positive network effects in information security and vulnerability disclosure.

Organisations can contract external security consultants to perform security tests and audits. External contractors are typically bound by the agreed budget, scope, time and location of the assignment, which can ultimately limit the possible vulnerabilities that can be identified. By using a CVD or bug bounty programme, organisations can receive an ongoing feed of vulnerability reports and, in the case of a bug bounty programme, only pay for valid vulnerabilities according to their severity (i.e. paying through a 'results-based', rather than a 'time-spent' based model).[120]

In addition to the perceived security benefits of vulnerability disclosure, organisations may also engage in CVD for any combination of three broader categories of motivations and incentives:

- **To raise awareness and engage with the security community.** CVD or bug bounty programmes can help raise awareness on information security matters within an organisation, particularly at the management level, as well as help raise awareness of the importance of information security more broadly in society.[121] Increased activity in CVD and bug bounty programmes have, in recent years, also resulted in significant attention on information security matters from media, policy makers and the wider public. Raising the importance of information security may help offset some of the perverse incentives to prioritise costs or speed to market over incentives (as discussed in Chapter 3). CVD and bug bounty programmes can also facilitate organisations to more productively engage with the security community by enabling a process or platform for two-way discussions, which can help to build trust and foster ecosystem thinking.[122]
- **In response to customer demand,** where the vendor develops a vulnerability disclosure process or programme to either strengthen overall product security or to showcase it as a proxy measure of comprehensive security practices. Within the current security environment, users – or particular user segments – may be more inclined to buy or engage with a vendor that is perceived to be investing in security or showcasing leadership in vulnerability disclosure. Vendor awareness of security demand therefore highlights the possibility of influencing vendor behaviour through customer or user preferences and demands.[123] If customers get increasingly security savvy and reduce the presence of information asymmetries, vendors may be more incentivised to further invest in security measures and CVD.

[118] Interview 3.
[119] Interviews 2, 4, 8, 9, 10, 13.
[120] Interview 7.
[121] Interview 4.
[122] Interviews 2, 9, 10.
[123] Interviews 2, 8, 9.

- **For ethical or social responsibility reasons** where vendors engage in vulnerability disclosure because they believe it to be an ethical or social responsibility to contribute to the overall strength of security or contribute to overall social welfare. This can, for example, be seen in 'The Internet Bug Bounty' programme, which is a bug bounty programme for core internet infrastructure and free open source software run by an independent and unpaid panel of security experts from the community and sponsored by Facebook, GitHub, Ford Foundation, Microsoft and HackerOne.[124]

However, not all organisations have a CVD policy in place or operate a vulnerability disclosure programme, and there are a number of barriers or disincentives that currently reduce the likelihood of an organisation's participation in vulnerability disclosure, including:

- The costs of implementation and operation of CVD or bug bounty programmes.
- A lack of awareness or understanding of vulnerability disclosure and how it could benefit the organisation.
- A lack of management support.
- A Lack of organisational or technical capacity.
- Legal barriers or uncertainty.

Costs of implementation and operation of CVD or bug bounty programmes can also disincentivise organisations to engage with vulnerability disclosure. Some organisations may perceive skewed cost-benefit calculations, where investment into security at large or CVD in particular is not an economically sound business decision that would yield sufficient returns on investment.[125] Organisations may also perceive CVD or bug bounty programmes as too costly, as organisations need to develop processes, policies and procedures for vulnerability disclosure and dedicate resources to the management and operations of disclosure programmes. The required resources will differ depending on the size of the organisation and the scope and nature of the vulnerability disclosure programme implemented, as well as the organisational and technical capacity of the organisation.

These cost concerns may also stem from concerns of additional work due to large volumes of invalid reports or noise, which may take valuable analyst time away from other, more important, security tasks. Previous research from the University of California, Berkeley and Pennsylvania State University has shown that bug bounty programmes may suffer from high error rates (i.e. large volumes of false positives), which may be considered invalid for a number of reasons (e.g. not a valid vulnerability, errors in the vulnerability submission, out of scope for the vulnerability programme etc.). In fact, several bug-bounty platforms have acknowledged that one of the primary challenges in operating a bug bounty programme is managing the volume of erroneous or invalid reports.[126] One study found that the volume of invalid reports can range between 35 to 55 per cent, which has significant resource implications.[127] However, careful consideration of the design of the vulnerability programme, as well as clear guidelines and rules, may help reduce the noise. Offering financial incentives may also help, as previous research from Pennsylvania State University has found that the average monetary bounty is positively correlated with the volume of valid reports received.[128]

[124] See https://internetbugbounty.org/ (As of 31 October 2018).
[125] Interview 2.
[126] Laszka et al. (2016).
[127] Zhao et al. (2017).
[128] Zhao et al. (2015).

Awareness of vulnerability disclosure has increased in recent years, but it is not yet standard practice in most business sectors and lack of management support may hamper CVD adoption.[129] Particularly in sectors where vulnerability disclosure is less common or among organisations that have less mature information security arrangements, it can be difficult for organisations to appreciate the potential security and economic benefits that CVD or bug bounty programmes could realise. Within this context, organisations may also be faced with a 'first mover' challenge where an organisation is reluctant to be the first organisation in their sector or particular context to implement CVD or bug bounty policy or programme. However, one interviewee pointed out that there are currently few sectors where there is not an industry leader with a progressive view of CVD.[130]

Lack of management support can also hamper an organisation's ability to implement a CVD or bug policy or programme.[131] A lack of understanding of and support for information security at the management level is a well-documented challenge, and this also extends to issues of vulnerability disclosure. This can be attributed to a lack of awareness or understanding of the issues at hand or a reluctance to internalise some of the externalities present in the information security market (i.e. letting other actors bear the cost of security). It can also be related to a distrustful view of the security community and a reluctance to share information about potential vulnerabilities. Members of the security community, or 'hackers', are oftentimes characterised as malicious actors that cannot be trusted, which has implications for vulnerability disclosure in that organisations may think of CVD as 'letting the bad guys in'.[132] Information asymmetries and adverse selection at the cost of security therefore also exists within organisations, typically between the information security department and senior management.

Organisations of lesser information security maturity may also be reluctant to share information about vulnerabilities in their products or services in fear of reputational damage or attacks.[133] Yet, as noted in Chapter 3, vulnerabilities are ubiquitous so, rather than opt for secrecy, organisations should recognise that all organisations will be faced with vulnerabilities and what ultimately matters is the ability to receive and respond to them.

Lack of organisational or technical capacity to design, implement and operate a CVD or bug bounty policy or programme may hamper an organisation's vulnerability disclosure work – even in the presence of management support and funding.[134] Designing and appropriately scoping a vulnerability disclosure policy can be one of the primary challenges faced by organisations, particularly if the organisation is a new entrant to vulnerability disclosure. These challenges may also extend to particular and specific logistical or operational barriers (e.g. How do can you remunerate or pay a security research that is outside your country? How do you define a legal agreement that provides a safeguard to researchers but does not jeopardise business interests?, etc.).[135] Organisational and technical capacity is also the enabler for operating a CVD or bug bounty programme. Organisations must have sufficient personnel, technical knowledge and capacity to receive, triage and develop remediation measures for vulnerabilities while maintaining efficient and transparent communications with security researchers.[136]

[129] Interviews 2, 8.
[130] Interview 8.
[131] Interviews 2, 8.
[132] Interviews 2, 4. 8.
[133] Interviews 2, 8, 10, 12.
[134] Interviews 4, 8, 12, 13.
[135] Interviews 4, 8.
[136] Interviews 1, 9, 10.

Legal barriers or uncertainty may also present themselves when implementing a vulnerability disclosure programme, particularly in relation to bug bounty programmes.[137] CVD or bug bounties can involve inviting largely unknown security researchers from anywhere in the world to explore and test an organisation's systems, which could have unintended consequences. Organisations could also be concerned about the behaviour of security researchers, who may jeopardise system integrity, collect commercially sensitive information and intellectual property, or disclose data or vulnerabilities to third-parties, competitors or the public.[138] It can also be challenging for organisations to navigate a complex legal landscape, particularly if operating in multiple national jurisdictions or dealing with platforms or security researchers in multiple countries. As seen in Section 4.3.2, organisations often face challenges when aligning CVD or bug bounty policies with End User License Agreements (EULAs) or other legal agreements.

Further to receiving vulnerability reports, vendors also p**lay a role in identifying remediation measures or patching identified vulnerabilities**. Previous research by Ashish Arora and Rahul Telang has shown that, due to the nature of the economic factors of the information security market (as discussed in Chapter 2), vendors tend to engage in less than socially optimal patching behaviour if left unregulated.[139] This is primarily driven by externalities in the information security market, where costs incurred by the exploitation of vulnerabilities are not borne by the vendors ultimately responsible for the vulnerability, as well as the prevalence of liability dumping or shifting between different actors across the supply chains. Patching behaviour is also influenced based on the nature of the vulnerabilities (e.g. higher impact vulnerabilities are patched more quickly), software type (e.g. open source software is generally patched faster than propriety software), and the type of update through which the vulnerability is patched (e.g. security patches are more quickly deployed than feature updates).[140] Vulnerabilities that affect multiple vendors are typically also patched more quickly than single vendor vulnerabilities.[141]

However, research from the University of North Carolina at Charlotte has also shown that patching behaviour can be influenced by regulations and vulnerability disclosure programmes in order to promote more socially desirable outcomes.[142] In general, the disclosure of vulnerabilities accelerates remediation work and release of appropriate patches. Vulnerability disclosure through a trusted partner, such as CERT/CC, may further speed up patching behaviour, highlighting the importance of the reputation of the disclosure coordinator.[143]

Within CVD or bug bounty programmes, organisations also have to make decisions of **whether or not to publish a particular vulnerability publicly** after it has been remediated. Organisations may want to publish information about the vulnerability to:

- Showcase the competence, communication and vulnerability disclosure good practice that the organisation performs to.
- Give back to the community, share lessons identified and incentivise further participation in vulnerability disclosure.

[137] Laszka et al. (2016).
[138] Zhao et al. (2017).
[139] Arora & Telang (2005).
[140] Temizkan et al. (2012).
[141] Arora et al. (2010).
[142] Temizkan et. al. (2012).
[143] Arora et al. (2010).

- Showcase accountability (i.e. that organisation is a responsible security actor that listens to the community and remediates valid vulnerability reports) and build trust in the vulnerability disclosure landscape (as to avoid an information security 'tragedy of the commons'.[144]

Organisations may however also choose not to disclose reported vulnerabilities publicly due to perceived security risks associated with public disclosure (e.g. that disclosure could reveal sensitive business practices or technical details that could enable other types of attacks). The ability to disclose vulnerabilities publicly may also be hampered by the regulatory or legislative environment of the organisation (e.g. companies operating in the financial sector).[145] Decisions are to not disclose remediated vulnerabilities publicly are perceived to be more common in private bug bounty programmes compared to other forms of vulnerability disclosure (see also discussion in Section 4.5.1).

4.4.2 Governments may take on several different roles in the vulnerability disclosure process, which affects its incentives and considerations

Governments are multifaceted actors in the vulnerability disclosure process and may perform several roles, including the role of finder, vendor/vulnerability owner, vulnerability coordinator, or those responsible for vulnerability stockpiling.

Governments may engage in vulnerability disclosure in a number of ways:

- **As a vulnerability coordinator or programme owner**, where a trusted government entity operates a responsible disclosure programme in which researchers can report vulnerabilities identified in government applications, networks of services.[146] Between 2017 and 2018, the Centre for European Policy Studies (CEPS) coordinated a Task Force on Software Vulnerability Disclosure in Europe that conducted work to help define guidelines to harmonise CVD processes in Europe.[147] The CEPS Task Force found that only three European countries have an established CVD process in place at the national level[148]; however, a number of countries are in the process of establishing a CVD policy.[149] In addition to operating a national CVD policy, governments can also sponsor semi-independent organisations to perform CVD, either on behalf of the government or for wider society.[150] Government organisations may also run vulnerability disclosure or bug bounty programmes in order to safeguard its own applications, networks or services.[151]
- **As a vulnerability finder**, where government organisations engage in security research or security testing and identify vulnerabilities, which can be done through regular and ongoing information security work or through dedicated efforts to identify vulnerabilities for national security purposes (see discussion below).

[144] Interviews 1, 5, 10.

[145] Interview 10.

[146] For example, the Netherlands National Cyber Security Centre's 'Responsible Disclosure' programme. See https://www.ncsc.nl/english/security (As of 31 October 2018).

[147] CEPS (2018).

[148] France, the Netherlands and Lithuania.

[149] Including Austria, Belgium, Bulgaria, Czech Republic, Finland, Germany, Hungary, Italy, Latvia, Luxembourg, Romania, Slovenia and the United Kingdom.

[150] For example, the US CERT Coordination Center (CERT/CC), which is part of the Software Engineering Institute (SEI), a not-for-profit federally funded research and development centre (a public-private partnership between Carnegie Mellon University and the US government).

[151] For example, the US Department of Defense 'Hack the Pentagon' programme. See https://www.hackerone.com/resources/hack-the-pentagon (As of 31 October 2018).

- **As a vulnerability buyer or keeper**, where a government actor within defence, national security or the intelligence services identifies or receives vulnerabilities that are subsequently not publicly disclosed or else are subject to delayed disclosure due to national security interests.

One of the key issues in recent years relates to governments engaging in vulnerability research and/or non-disclosure of vulnerabilities for national security purposes, particularly for intelligence activities or for offensive cyber capabilities.[152] On the one hand, if a government finds or receives information about a vulnerability, it can choose to keep it secret and, as such, theoretically gain a competitive military, intelligence or economic advantage over its adversaries (who may not yet know about the vulnerability). The equity may enable the government to execute its mission more efficiently or achieve additional effects through its cyber capabilities.[153] On the other hand, if a government chooses not to disclose the vulnerability to vendors, it may have significant wider effects on overall security. This is particularly relevant if the vulnerability affects multiple vendors or if it has severe potential impacts on individuals, organisations or society, if exploited. An example of the possible impacts of non-disclosure can be seen in the EternalBlue case study featured in Chapter 5.[154]

Non-disclosure of vulnerabilities for national security purposes presents a particular challenge due to the nature of vulnerabilities, particularly the zero-day variety, which may incentivise stockpiling rather than responsible disclosure. Research from the RAND Corporation has shown that zero-day vulnerabilities enjoy a relatively long shelf-life (i.e. they can be actively exploited for an average of seven years) and a low probability of discovery by another party.[155] Previous research that used a game-theoretic model of state behaviour – in which states had to choose between mutual protection (i.e. responsible vulnerability disclosure) and offence at risk (i.e. keeping the vulnerability secret) – also illustrated that states were rewarded for pursuing offensive strategies.[156]

Part of this challenge stems from a lack of clarity or consistency as to how a government will decide whether a vulnerability is severe enough to warrant disclosure or not. The United States has implemented the ‘Vulnerabilities Equities Policy and Process’ (VEP) in order to facilitate this process and ensure adequate levels of transparency and oversight in relation to how these decisions are made.[157] There has also been a discussion about the feasibility of implementing similar disclosure review processes in Europe (what CEPS refer to as government disclosure decisions processes (GDDP)).[158]

Governments also influence the vulnerability disclosure landscape in their role as enactors and implementers of legislation, regulation and policy.[159] Governments can advocate and regulate organisations in order to implement CVD policies – either by mandating their use or by sharing information or good practice in order to build awareness. The findings of the CEPS Task Force highlight the important role that governments can play in terms of leading by example through the implementation of government

[152] Delcheva & Soesanto (2018).
[153] Ablon & Bogart (2017).
[154] EternalBlue is an exploit of a number of vulnerabilities allegedly discovered and kept secret by the US National Security Agency (NSA).
[155] Ablon & Bogart (2017).
[156] Moore et al. (2010).
[157] United States Government (2017).
[158] CEPS (2018).
[159] An examination of how legislation, regulation and policy may affect vulnerability disclosure is featured in Section 4.5.

CVD policies at the national level.[160] However, while government involvement in CVD is crucial, one of its primary tasks should be to explicitly take a step back and re-affirm that CVD is ultimately a process between security researchers and vendors (and if required, a coordinator).[161] This is also the approach adopted by the Dutch government in its CVD guidelines, which overtly states that the CVD process is primarily a matter for organisations and reporting parties (and not the government). However, the guidelines also make it clear that the National Cyber Security Centre can assist in a disclosure process if requested, particularly in sharing information on the vulnerability with its constituency in order to limit further security risks arising from the vulnerability.[162]

4.5 Structural considerations and incentives

In addition to the contextual economic factors of information security discussed in Chapter 2, the economics of vulnerability disclosure is also influenced by a number of structural factors. As noted in the previous sections of this chapter, behaviour in vulnerability disclosure can be influenced by individual or organisational considerations and incentives, which in turn can be affected by structural elements of the wider disclosure landscape. These structural considerations may include legislation related to vulnerability disclosure, regulation of vulnerability disclosure or vulnerability markets, and liability mechanisms such as cyber insurance.[163]

4.5.1 Legislation plays a central role in the vulnerability disclosure landscape

In the context of this study, legislation refers typically to laws passed by a national, regional or supranational legislature. These laws seek to interpret the needs of modern information society and balance these with the need for security and liability, which could, for example, relate to restricting vulnerability information, limiting liability for software or hardware errors, preventing the reverse engineering of software, or exporting sensitive technologies.[164] There is typically no legislation specific to vulnerability disclosure, but there are many types of legislation that have direct effects on behaviour in the vulnerability disclosure process, particularly legislation related to unauthorised access to computer systems and networks.[165] Legislation that clearly outlines what is legally admissible to do as part of a vulnerability identification process may put security researchers at ease, allowing them to undertake vulnerability finding activities without risking legal action. In contrast, overly restrictive or misaligned legislation may, in turn, have chilling effects on vulnerability disclosure and result in fewer security researchers engaging in vulnerability identification and disclosure.[166]

Given the risks of malicious actors exploiting undiscovered vulnerabilities in European applications, networks and services, there may be a need for an appropriate legal framework that enables the discovery of vulnerabilities under explicitly agreed circumstances. As noted in Section 4.3.2, security researchers often quote legal jeopardy as one of the reasons not to engage in vulnerability disclosure, and many current legal solutions – including contracts between security researchers and programme owners – may put unnecessary legal risk on the individual researchers.

The CEPS Task Force report on CVD features an extensive discussion of how vulnerability disclosure is addressed in a number of European jurisdictions, as well as at the European Union level (including criminal

[160] CEPS (2018).
[161] Interviews 5, 8, 9.
[162] The Netherlands National Cyber Security Centre (2018).
[163] Granick (2005); Böhme (2005); Camp & Wolfram (2004); Böhme & Kataria (2006).
[164] Takanen et al. (2004).
[165] Often referred to as ‘hacking laws’.
[166] Vaas (2018).

law, data protection law, and other regulatory issues such as copyright, trade secrets, patents, trademarks and export control regulation).[167] The recommendations put forward in the CEPS Task Force include to amend Directive 2013/40/EU on attacks against information systems (the 'EU cybercrime Directive'), to support CVD, as well as to consider CVD in the implementation of the Directive on security of network information systems (NIS Directive), the General Data Protection Regulation (GDPR) and other national legislative and non-legislative activities.[168] The CEPS Task Force also recognises the importance of legal safeguards for researchers (safe harbour) to ensure that the legal liability and responsibilities of security researchers are clarified to enable them to perform their work without fear of prosecution.[169]

One of the commonly discussed issues concerning legal regimes for CVD is whether or not to make CVD policies mandatory for vendors. On the one hand, a legal requirement would increase vendor uptake of CVD policies and facilitate the reporting and remediation of discovered vulnerabilities. On the other hand, mandatory CVD policies could also result in adverse effects on the vulnerability disclosure ecosystem. As highlighted in Section 4.4, effective CVD requires appropriate levels of organisational and technical maturity within vendors to be able to receive, understand and respond to reported possible vulnerabilities. The adoption of CVD policies by less mature organisations due to legal requirements could therefore lead to negative disclosure experiences where vulnerabilities are reported but not followed up or resolved, or finders are met by uncommunicative or hostile vendors, which may lead to researchers being discouraged to participate in CVD processes.[170]

4.5.2 Regulation and liability can provide structure to the vulnerability disclosure process and incentivise desired behaviours

Regulation can be issued at the governmental or organisational level, for example by an organisation that operates a vulnerability market. Regulation exists in several forms, including market regulation and liability regimes.

Proponents of regulated vulnerability markets argue that such markets may bring several benefits, including increased incentives for security researchers to identify and responsibly disclose vulnerabilities, as well as increased opportunities to patch and remediate vulnerabilities before they are made public.[171] One study has also found that the use of regulated vulnerability markets delays the onset and reduces the impact of attacks utilising a given vulnerability, and decreases the risk of first attack and the overall volume of attacks using that vulnerability.[172]

However, vulnerability markets, particularly unregulated ones, may also result in negative incentives or outcomes, including that:

- More identified and disclosed vulnerabilities may lead to less security and an increase in security incidents and attacks.[173]
- The introduction of private sector intermediaries in the vulnerability disclosure process may result in additional complexity or loss of trust between finders and vendors.[174]

[167] See Chapter 4 in CEPS (2018).
[168] See chapter 6 in CEPS (2018).
[169] CEPS (2018).
[170] Interview 9.
[171] Ransbotham et al. (2012).
[172] Ransbotham et al. (2012).
[173] Mullin (2001).
[174] Li & Rao (2007).

- Actors may be more incentivised to develop products or services with vulnerabilities in order to sell these in vulnerability markets or be more incentivised to leak information to make their services or markets more valuable.[175]

The existence of different types of vulnerability markets therefore influences behaviour within the vulnerability disclosure process and may incentivise different actors to act in a certain way. Regulated markets may have more transparent results, particularly in relation to socially beneficial outcomes and positive incentives, but further research evaluating their performance and effects is required.

Ensuring the security of increasingly connected and complex systems will require vendors and operators to engage in holistic and thorough security practices, of which vulnerability disclosure plays an important part. Regulators and governments should consider if vendors will enact these practices without intervention or if additional incentives are required. **Liability regimes** can take many different forms but typically refer to legal or regulatory measures to hold vendors or other parties that make products or services available in a market accountable for costs or damages incurred by those products or services. Liability incentives can, as such, help incentivise more secure systems by offering reduced liability for compliant systems or increased liability levels for insecure systems (or systems that do not comply with CVD good practice), thus reducing the presence of externalities in the information security market.

If vendors are held financially liable for damages incurred by vulnerabilities in their products and services, the vendor would, in theory, be more incentivised to proactively engage in CVD or bug bounty programmes in order to identify and remediate vulnerabilities before they are exploited. In other words, liability regimes could help correct some of the externalities present in the information security market, as discussed in Chapter 3. However, effective use of liability regimes and mitigation of moral hazards require clear connections between the vulnerability and the measurable harm caused by its exploitation, which are not always straightforward to make especially in a complex supply chain. In theory, liability regimes could also improve the speed and quality of patches, as vendors would seek to avoid the exploitation of vulnerabilities.[176] Liability regimes may become increasingly important in relation to cyber-physical and IoT systems that will be in operation, and therefore will have to be kept secure, for long periods of time. If there is a third-party system embedded in a car that is produced and sold today, it is unclear who would be responsible for identifying, reporting and resolving vulnerabilities in that system for the next 20 to 30 years.[177]

In the context of reporting vulnerabilities, liability can also be extended to the finder or reporter of the vulnerability. As discussed in Section 4.3.2, it is important to ensure legal safeguards for security researchers participating in CVD to prevent them from being blamed or held liable for vulnerabilities or their exploitation. There is pre-existing EU regulation that could be used as inspiration for CVD specific legal and regulatory safeguards for CVD and bug bounty programmes. The EU regulation on reporting, analysis and follow-up of occurrences in civil aviation seeks to prevent organisations to use reported vulnerability information against the reporter without granting immunity to reporters in case of gross negligence, willful violations and destructive acts. Similar regulation could feasibly be used to ensure legal safeguards for CVD and to promote participation in CVD and bug bounty programmes.

Box 1 EU Regulation on reporting of occurrences in civil aviation: Article 15: Confidentiality and appropriate use of information

[175] Kannan & Telang (2005).

[176] Cavusoglu et al. (2005).

[177] See Leverett, Clayton & Anderson (2017) for an extensive examination of these issues.

> 2) Without prejudice to the provisions relating to the protection of safety information in Articles 12, 14 and 15 of Regulation (EU) No 996/2010, information derived from occurrence reports shall be used only for the purpose for which it has been collected. Member States, the Agency and organisations shall not make available or use the information on occurrences:
>
> (a) in order to attribute blame or liability; or
>
> (b) for any purpose other than the maintenance or improvement of aviation safety.
>
> 4) Member States shall ensure that their competent authorities referred to in Article 6(3) and their competent authorities for the administration of justice cooperate with each other through advance administrative arrangements. These advance administrative arrangements shall seek to ensure the correct balance between the need for proper administration of justice, on the one hand, and the necessary continued availability of safety information, on the other.

Source: Adapted from Peeters (2017)/Regulation (EU) 376/2014 on the reporting, analysis and follow-up of occurrences in civil aviation.

4.5.3 Insurance represents another type of structural incentive that can influence behaviour

Another structural lever that can influence the vulnerability disclosure process is **insurance**. Insurance is typically used to address issues that cannot be reasonably mitigated by security measures due to its complexity or associated costs. Cyber insurance is a relatively new application of this principle and typically refers to insurance policies designed to address first- and third-party losses due to cyber attacks or incidents (due to malicious actions, malfunction or errors).[178]

Despite substantial growth in the cyber insurance market,[179] much of the academic literature in this field remains theoretical in nature.[180] While empirical analysis of the effects of cyber insurance is lacking, the theoretical literature illustrates that the level of provision of cyber insurance very much interacts with other economic factors of information security. One framework shows that cyber insurance comprises five factors: the networked environment (e.g. network effects and externalities), supply and demand-side factors (e.g. risk aversion, utility functions of firms, and the competitive insurance landscape), the information structure (e.g. adverse selection and moral hazards), and the organisational environment (e.g. regulatory incentives).[181]

In addition to its interaction with the economic factors of information security, cyber insurance can also influence behaviour in the vulnerability disclosure process. Cyber insurance could contribute to either:

- **Improved vulnerability disclosure (and security)**, if the development and adoption of structured vulnerability disclosure processes or programmes are mandated by insurance companies and placed as a requirement for coverage, positive social benefits can be expected. This is particularly true in the case of continued growth of the cyber insurance market and the utilisation of cyber insurance as one part of a proactive security strategy by firms.
- **Reduced overall security**. Risk management is often perceived as a trade-off between investing in sufficient security controls to reduce the average loss of a security incident and insuring against the loss. This means that if cyber insurance is cheap, firms may be less incentivised to provide

[178] Romanosky et al. (2017).
[179] ENISA (2016).
[180] See, for example: Böhme & Gaurav (2006); Baer & Parkinson (2007); Böhme (2010), Johnson et al. (2011); Böhme & Schwartz (2010).
[181] Böhme & Schwartz (2010).

adequate information security controls to protect against losses.[182] Firms would therefore indirectly place the cost of security on the insurers, rather than bearing it themselves, which could reduce the overall security of the information security market. Equally, if insurance is costly, firms may be incentivised to further invest in information security controls rather than to rely on insurance or to drive down the cost of their premiums.[183] The price and coverage of cyber insurance may therefore have direct implications for the vulnerability disclosure process.

However, research into cyber insurance and vulnerability disclosure is limited. Further empirical studies are needed to better understand the effects that cyber insurance has on vulnerability disclosure and the ways in which insurance may influence the behaviour of different actors in a vulnerability disclosure process.

4.6 Normative considerations and incentives

Normative considerations such as **voluntary standards, good practice and norms** can also influence behaviour in the vulnerability disclosure process. As previously noted, **individual ethics or expectations of behaviour** may guide finders in the vulnerability process and, similarly, industry ethics or norms may act as structural incentives for the behaviour of other agents in the vulnerability disclosure process. Professional codes of conduct such as the IEEE Code of Ethics,[184] and corporate initiatives such as the Cybersecurity Tech Accord,[185] which promote responsible behaviour in cyberspace, may incentivise positive behaviour across the information security market.

Another element that can help incentive positive behaviour across the vulnerability disclosure landscape is the development and implementation of **good practice, standardised processes and standards**. These factors are not typically considered the primary drivers of behaviour, they may nevertheless facilitate shared understanding, contribute to capacity building and highlight the importance of responsible vulnerability disclosure to a wider audience of relevant actors. Standards such as 'ISO/IEC 29147: Vulnerability disclosure' and 'ISO/IEC 30111: Vulnerability handling processes' provide authoritative reference documents for vulnerability disclosure, while good practice documents from reputable actors such as the CERT Coordination Center (CERT/CC) 'CERT Guide to Coordinated Vulnerability Disclosure' and the Forum of Incident Response and Security Teams (FIRST) 'Guidelines and Practices for Multi-Party Vulnerability Coordination and Disclosure'[186] may help set expectations about behaviour in the vulnerability disclosure process. Standardisation of vulnerability disclosure language and adoption of voluntary standards may also help reduce the information burden on researchers and organisations, increase awareness and reduce transaction costs.

Additionally, recognition of the importance of vulnerability disclosure processes in other information security-related standards or frameworks may further incentivise organisations to implement and operate vulnerability disclosure processes. For example, the inclusion of vulnerability disclosure processes in the US National Institute of Standards and Technology (NIST) Framework for Improving Critical Infrastructure Cybersecurity may influence the uptake and use of vulnerability disclosure across US federal authorities.[187] In Europe, a number of initiatives have also highlighted the important role of vulnerability disclosure, including the recognition from the European Council who 'welcome[d] the call to acknowledge the

[182] Kesan et al. (2005).
[183] Romanosky et al. (2017).
[184] See https://www.ieee.org/about/compliance.html (As of 31 October 2018).
[185] See https://cybertechaccord.org/ (As of 31 October 2018).
[186] FIRST (2017).
[187] NIST (2018).

important role of third party security researchers in discovering vulnerabilities in existing products and services and call[ed] upon Member States to share best practices for coordinated vulnerability disclosure'.[188] Active support to CVD from European Institutions can also be seen in the European Commission's EU-Free and Open Source Software Auditing Community (EU-FOSSA) programme, which includes a bug bounty programme for vulnerabilities in open source software.[189] Government and other institutions can therefore contribute to awareness and profile-raising for CVD issues, as well as act as advocates for not-for-profit and other CVD actors.

4.7 Emerging trends in vulnerability disclosure

Vulnerability disclosure actors may adjust their behaviour in accordance with the type of vulnerability disclosure process that is undertaken; a CVD process and a bug bounty disclosure process may not be realised in the same way. In recent years, the prevalence and participation in bug bounty programmes have increased dramatically. Since their inception during 2012–2013, BugCrowd and HackerOne have together attracted thousands of organisations, tens of thousands of security researchers, have responded to over hundred thousand vulnerabilities and paid out tens of millions in bug bounties.[190]

The growth in bug bounty programmes and the bug bounty economy has impacted the vulnerability disclosure landscape in a number of ways. The increased use of bug bounty programmes has helped to raise awareness of vulnerability disclosure and perhaps contributed to a 'normalisation' of vulnerability disclosure – where an increasing number of organisations pro-actively engage with the security community. It has also increased and perhaps improved the interaction between the security community and vendors, fostering a mutual understanding between the two groups.[191] It is also evident that the growth of bug bounty programmes has resulted in more money being paid to security researchers for identifying and responsibly reporting vulnerabilities.[192] The increased monetary presence of these types of programmes may also work to inspire a new generation of information security professionals, which could help address the prominent shortage of suitably qualified and experienced professionals globally.[193]

Bug bounty pay-outs can also help improve the understanding of the dynamics between regulated and black markets, particularly in relation to prices of vulnerability information or exploits. Whereas certain bug bounty programmes pay out significant bounties (>US$10,000), the average bounty is considerably lower. Some platforms, such as Bugcrowd, provide practical guidance on how to approach vulnerability and bounty pricing. The 'Defensive Vulnerability Pricing Model' is based on two aspects: organisational maturity and vulnerability priority (i.e. its technical and business impact). There are three levels of organisational maturity (Basic, Progressing and Advanced) that interrelate with five vulnerability severity levels, as illustrated in Table 4.1.[194]

Table 4.1 Bugcrowd's Defensive Vulnerability Pricing Model

[188] European Council (2017).

[189] EU-FOSSA is managed by the European Commission's Directorate General of Informatics (DIGIT) and implements the European Parliament's Pilot Project 'Governance and quality of software code – Auditing of free and open source software'. See https://joinup.ec.europa.eu/collection/eu-fossa-2 (As of 31 October 2018).

[190] Bugcrowd (2018); HackerOne (2018).

[191] Interviews 2, 9, 10.

[192] Interviews 1, 2, 9.

[193] The 2017 Global Information Security Workforce Study projects a shortage of 1.8 m information security professionals globally by 2022 (GISWS 2017).

[194] The priority levels are: P1 – Critical, P2 – High, P3 – Medium, P4 – Low, and P5 – Acceptable risk. Further definitions and example vulnerabilities for each of the priority levels can be found in Bugcrowd (n.d.).

	BASIC	PROGRESSING	ADVANCED
Pay-out range	US$100–1,500	US$200–5,000	US$300–15,000
Average bug pay-out	$300	$600	$1000
P1	$1,500	$5,000	$15,000
P2	$900	$1,800	$2,500
P3	$300	$600	$900
P4	$100	$200	$300

Source: Bugcrowd (n.d.).

In addition to this type of practical guidance, organisations may also examine comparative organisations to understand vulnerability bounty levels.[195] This particularly applies to mature organisations that wish to market their bug bounty programmes in order to attract and retain top security researcher talent.[196]

Compared to grey or black markets, these bug bounty levels may seem low. Previous research from the RAND Corporation has shown that prominent zero-day vulnerability exploits can be sold for around US$30,000–50,000 on the black market, and for as much as US$50,000–300,000 in the grey or government markets.[197] Beyond zero-day exploits, general exploit kits are typically available on the black market for a few thousand USD.[198] However, developed exploits, particularly for zero-day vulnerabilities, will always be more costly than comparable vulnerability information disclosed through bug bounty programmes. In general, for any vulnerability market, the associated cost or reward level is assessed in relation to a number of factors: anticipated impact, ease of discovery, and frequency of vulnerabilities in the particular product.[199] The main differentiator between defensive cyber security and offensive cyber security activities, pricing is in essence that for the first case the aim is to effectively kill the vulnerability (i.e. to develop and implement an appropriate remediation measure), whereas in the second case, the aim is to actively exploit the vulnerability (thereby generating mission and lifecycle value).[200]

Lastly, bug bounty programmes are actively giving back to the community in terms of information sharing, capacity building and training.[201] These types of activities can help improve the security researcher community over time and also assist newer researchers in meeting the requirements of modern CVD (e.g. communication skills), as well as contribute to the wellbeing of the wider information security ecosystem.

However, there are also several concerns about how the prominence of bug bounty programmes may affect the information security and vulnerability disclosure environments. Bug bounty programmes may have limited practical security impact as most programmes are aimed at solving seemingly low-level or

[195] Interview 10.
[196] Bugcrowd (n.d.).
[197] Ablon & Bogart (2017).
[198] Ablon et al. (2014).
[199] I.e. high-impact vulnerabilities that required significant resources to discover, particularly in products with few known vulnerabilities will garner a high valuation; see Ablon & Bogart (2017).
[200] Interview 2.
[201] Interviews 2, 8.

common vulnerabilities (e.g. cross-site scripting vulnerabilities[202] or SQL injections[203]) in non-critical systems (e.g. public-facing websites, web applications, etc.).[204] One of the underlying rationales for bug bounty programmes (i.e. to attract members of the wider security researcher community to identify vulnerabilities) also presents practical limitations to which applications, systems and services that can be made available for testing in a semi-public setting.[205] This means that a significant portion of bug bounty activity is spent solving already well-known and perhaps not so critical vulnerabilities, rather than solving more serious and structural issues in the current ICT ecosystem.[206]

An increased uptake of bug bounty programmes may also have spillover effects on the wider information security research environment, as bug bounty programmes only pay out bounties to the finders (and not to other researchers who may have enabled the discovery). [207] Bug bounty platforms are also economic actors in their own right, and that may enact less than socially optimal decisions driven by their own business incentives. [208] Most significantly, bug bounties may result in less participation in non-paid CVD programmes if researchers are increasingly demanding or expecting monetary rewards for reporting security vulnerabilities or if bug bounties are predominantly operated as private programmes.[209] In a worst-case scenario, private bug bounty programmes could impose overly strict vendor terms and non-disclosure agreements that gives vendors complete control of the disclosure process, which could prevent further public or coordinated disclosure (even if a vulnerability is present elsewhere or have wider supply-chain implications).[210]

In the context of the European Union, it is also worth noting that the most prominent bug bounty platforms, which attract the most organisations and the top research talent, are based in the United States and may therefore provide limited value to Europe. There may be economic, legal, and other benefits to CVD adoption in Europe if there were prominent European alternatives to these US platforms.[211] This would enable European security researchers to participate in European CVD programmes and also help improve the European information security talent pool.

[202] Cross-site scripting, often abbreviated as XSS, is a type of vulnerability generally found in web applications that enables attackers to inject client-side scripts into web pages accessed by other users. See https://www.owasp.org/index.php/Cross-site_Scripting_(XSS) (As of 31 October 2018).
[203] SQL injections allow an attacker to introduce (or 'inject') code into a vulnerable SQL data-management applications and change the course of execution. A successful SQL injection exploit could enable an attacker to read sensitive data from a database, modify database data or execute administration operations. See https://www.owasp.org/index.php/SQL_Injection (As of 31 October 2018).
[204] Interview 3.
[205] Some of these may be circumvented through the use of private bug bounty programmes but social good of private bug bounty programmes may also be marginal.
[206] For example, over 70% of surveyed researchers active on HackerOne stated that their preferred product or platform to hack is websites. Some 28% of researcher preferred searching for XSS vulnerabilities, followed by SQL injection (23.1%) (HackerOne (2018)). Similarly, 57% of all programmes launched on BugCrowd in 2017/2018 concerned website targets, over 80% of paid bounties were for website-associated vulnerabilities, and 13% of all bounties related to XSS, 12% to server side injection and 9% to broken authentication and session management.
[207] Interview 3.
[208] Interview 13.
[209] Interview 1, 4, 9, 12, 13.
[210] Interview 12.
[211] Interview 4. There are currently several smaller European bug bounty and vulnerability platforms in Europe including BountyFactory (https://bountyfactory.io) and Intigriti (https://www.intigriti.com) (As of 31 October 2018).

4.8 On reflection

Vulnerabilities in software and hardware are common, persistent and have the potential to cause significant societal harm. Vulnerability disclosure is a complex process that often involves multiple actors, particularly when a coordinated disclosure process is pursued or the vulnerability affects multiple vendors. Vulnerability disclosure also takes place in an information security market that is subject to unique dynamics and influencing factors, which often produce incentives or behaviour that reduces security or increases the presence of vulnerabilities. The information security market has direct effects on how actors behave in the vulnerability disclosure process and may shape both the vulnerability disclosure process and its potential outcomes.

It is clear that the behaviour of actors within vulnerability disclosure is influenced by individual and organisational economic considerations of incentives, which may be complementary or conflicting. An overview of how these factors feature in a CVD process is illustrated in Figure 4.3.

Figure 4.3 Economic incentives, motivations and barriers in a CVD process

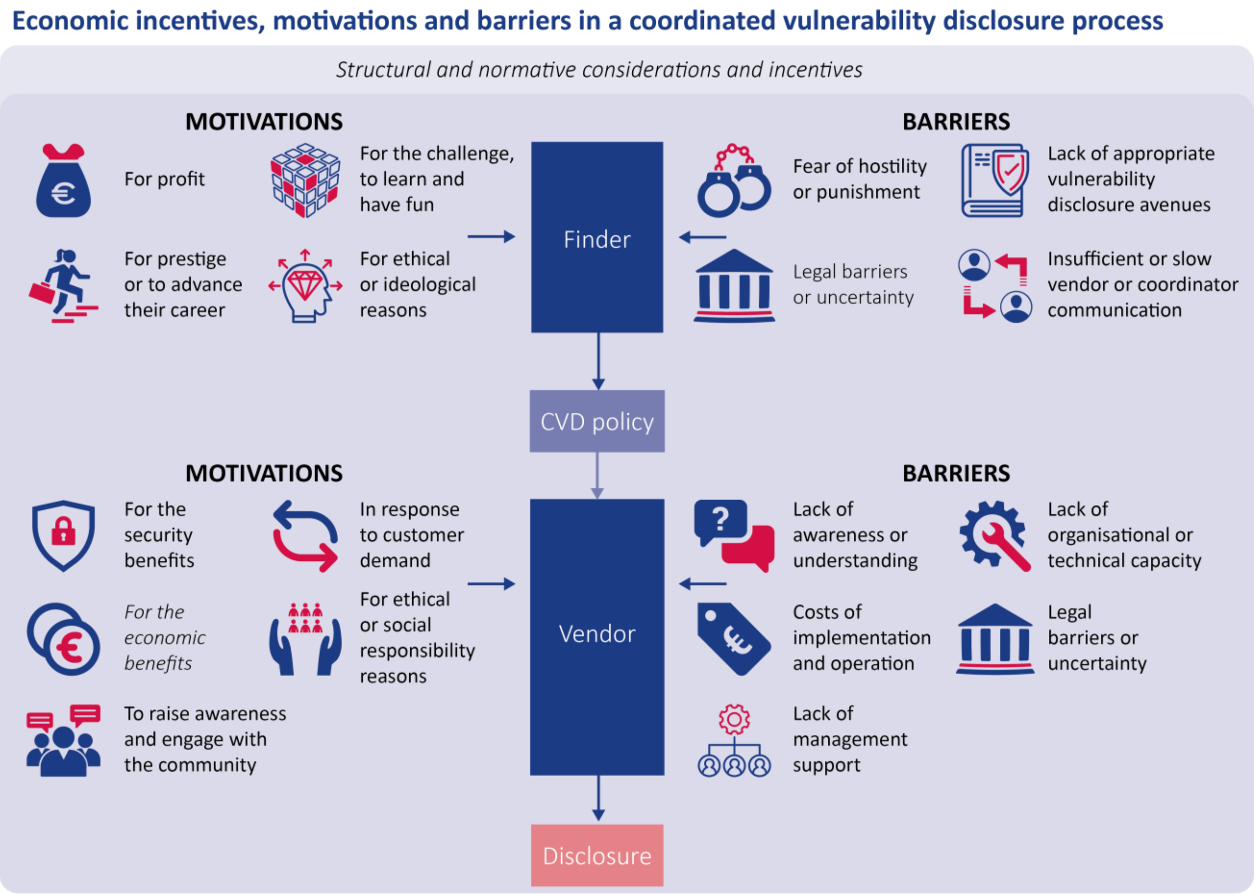
Economic incentives, motivations and barriers in a coordinated vulnerability disclosure process
Structural and normative considerations and incentives
MOTIVATIONS
For profit
For the challenge, to learn and have fun
For prestige or to advance their career
For ethical or ideological reasons
Finder
BARRIERS
Fear of hostility or punishment
Lack of appropriate vulnerability disclosure avenues
Legal barriers or uncertainty
Insufficient or slow vendor or coordinator communication
CVD policy
MOTIVATIONS
For the security benefits
In response to customer demand
For the economic benefits
For ethical or social responsibility reasons
To raise awareness and engage with the community
Vendor
BARRIERS
Lack of awareness or understanding
Lack of organisational or technical capacity
Costs of implementation and operation
Legal barriers or uncertainty
Lack of management support
Disclosure

Source: ENISA study on the economics of vulnerability disclosure

Some of these incentives may result in positive behaviour that produces net welfare gains in the information security market, whereas other incentives may drive negative or destructive behaviour. Both organisations and individuals also face barriers to participation in vulnerability disclosure, including barriers related to awareness, communication and legislation or regulation. While economic incentives and motivations may different between different actor groups, and within individual actor groups, vulnerability disclosure cannot take place without the presence of at least finders and vendors. The vulnerability disclosure landscape can therefore be conceptualised as an ecosystem that flourishes in the presence of mutually beneficial economic incentives and motivations.

However, even in the presence of undesirable incentives there may be opportunities to affect behaviour using different incentives or levers. Structural levers, such as legislation or regulation, can be important policy tools to influence the behaviour of different vulnerability disclosure actors to achieve socially desirable security outcomes. Legislation and regulation may help offset some of the negative consequences of the economic features of the information security market, as discussed in Chapter 3; increased liability for vendors may reduce externalities; better information and labelling of product security may reduce information asymmetries between users and vendors; and improved legal frameworks may reduce liability shifting between actors.

Chapter 3 and Chapter 4 have illustrated some of the relationships between the economics of the information security market, the economic incentives and motivations of vulnerability disclosure actors, and other factors that influence vulnerability disclosure including regulation, legislation and normative issues. The prevalence of economics in vulnerability disclosure emphasises the importance of a well-developed understanding of the economic aspects of vulnerability disclosure and how they may influence different processes. However, as noted in Chapter 4, there are currently research gaps in certain areas of the research field that may impede a full understanding of the economics of vulnerability disclosure. The identified research gaps are expanded upon in Section 6.2.

The preceding chapters have predominantly focused on academic or theoretical explorations of the economics of vulnerability disclosure, but it also necessary to empirically illustrate the economics of vulnerability disclosure. As such, Chapter 5 comprises two case studies of recently disclosed high-profile vulnerabilities and examines how their disclosure processes were carried out in practice, in addition to what economic considerations and incentives were present in those disclosure processes.

5. Vulnerability case studies

To illustrate how vulnerability disclosure takes place in practice and how economic considerations and incentives may influence it, this chapter presents two case studies of recently disclosed vulnerabilities: Meltdown/Spectre and EternalBlue. The case studies were chosen due to their prominence, as well as the fact that they represent both software and hardware vulnerabilities and are examples of responsible disclosure and non-disclosure.

For each of the two case studies, the analysis covers the nature of the vulnerability, the disclosure process and the aftermath and impact of the disclosure, particularly in relation to the economics of vulnerability disclosure.

5.1 Meltdown and Spectre

Meltdown and Spectre are two separate but closely related vulnerabilities discovered in 2017. In contrast to many recently disclosed and published high-profile vulnerabilities, Meltdown and Spectre were both vulnerabilities in hardware rather than software.

Modern processors work by performing what is known as speculative execution, which means that they try to guess ahead and execute commands before knowing whether these commands are correct or not. [212] By taking advantage of this process, the Meltdown and Spectre vulnerabilities are capable of causing the CPU to do speculative execution of code, while timing memory accesses to infer what has or has not been cached, to disclose the contents of memory.[213] Meltdown works by ‘melting’ the security mechanism in place, preventing random user programmes from accessing system memory.[214] Spectre misleads applications into purposefully opening various locations in their memory and is commonly perceived to be the more sinister of the two.[215]

5.1.1 Meltdown and Spectre were discovered by three security teams working independently of each other

Meltdown was independently discovered by three separate teams: Jann Horn from Google Project Zero; Werner Haas and Thomas Prescher from Cyberus Technology; and Daniel Gruss, Moritz Lipp, Stefan Mangard and Michael Schwarz from the Graz University of Technology. Spectre was independently discovered by two different groups: Jann Horn from Google Project Zero; and another group comprised of Paul Kocher, Daniel Genkin from the University of Pennsylvania and the University of Maryland, Mike Hamburg from Rambus, Moritz Lipp from Graz University of Technology and Yuval Yarom from the University of Adelaide and Data61. As such, all but one finder represented an academic institution (Jann Horn being from the private sector vulnerability research initiative ‘Project Zero’, a part of Google). The academic finders were in part funded by public research grants and other financial support, including from

[212] There are actually a total of five variants of the Meltdown/Spectre vulnerability: Variant 1: Bounds Check Bypass – CVE-2017-5753; Variant 2: Branch Target Injection – CVE-2017-5715; Variant 3: Rogue Data Cache Load – CVE-2017-5754; Variant 3a: Rogue System Register Read – CVE-2018-3640; and Variant 4: Speculative Store Bypass – CVE-2018-3639. For additional details, see https://www.us-cert.gov/ncas/alerts/TA18-004A (As of 31 October 2018).
[213] Bright (2018).
[214] Lipp et al. (2018).
[215] Kocher et al. (2018).

the European Research Council, the US National Science Foundation (NSF) and the US Defense Advanced Research Project Agency (DARPA).[216]

Surprisingly, the teams reported the vulnerability within months of one another – an amazing coincidence considering the chips have contained the flaw since the 1990s. By the time the Graz University researchers informed Intel in June 2017, the chip-makers had already been made aware of the issue on three separate recent occasions.[217]

5.1.2 The disclosure process highlights the complexity in multivendor CVD

Success in multivendor coordination requires both a technical understanding of the vulnerability at hand and an understanding of human communication and behaviour, as well as economic incentives and motivations at play.[218] This is particularly true if there is a technically complex vulnerability or challenging vendor or supply chain implication; or in the case of Meltdown and Spectre, a combination of both.

Complex supply chains make it more challenging to understand the impact the vulnerability may have across the supply chain and contribute to confusion as to who is ultimately responsible for coordinating, communicating and eventually remediating the vulnerability. Vulnerabilities can affect two types of supply chains: vertical and horizontal. In a vertical supply chain, a vulnerability may manifest itself across the supply chain as multiple products share a dependency on a vulnerable library or component. This would require the owner of the library or component issuing a patch and other actors in the supply chain implementing it for their products or services, which often leads to cascading effects where significant groups of users are left vulnerable while they await a patch from their particular product or service. In a horizontal supply chain, the same vulnerability could be found in multiple different products from different vendors due to vulnerabilities stemming from underspecified protocols, design flaws, etc. While these types of vulnerabilities are rare, they typically require significant resources to coordinate and mitigate, as multiple vendors have to develop patches for their implementations, as well as coordinate and mitigate effects in the downstream vertical supply chains.[219]

In the case of Meltdown and Spectre, there were both horizontal and vertical supply chain implications, and identifying optimal disclosure timing for all parties was challenging. The coordination process also highlighted the difficulty in deciding which vendors across the supply chains should be part of the coordination process, and be informed in an advance or simply notified at the point of disclosure. Google Project Zero typically invokes a 90-days disclosure period for identified vulnerabilities.[220] However, due to the technical complexity and perceived impact of the vulnerabilities, Project Zero invoked an 'extraordinary circumstances' clause for Meltdown and Spectre variants 1-3, effectively extending the disclosure period to several months.[221]

[216] In detail, the work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 681402), NSF awards #1514261 and #1652259, financial assistance award 70NANB15H328 from the US Department of Commerce, National Institute of Standards and Technology, the 2017–2018 Rothschild Postdoctoral Fellowship, and the Defense Advanced Research Project Agency (DARPA) under Contract #FA8650-16-C-7622.

[217] Greenberg (2018).

[218] In multivendor CVD, the technical details of the vulnerability may therefore dictate what the response should be, while social behaviour largely dictate how that response is realised. See Householder et al. (2017).

[219] Householder et al. (2017).

[220] See https://googleprojectzero.blogspot.com/2015/02/feedback-and-data-driven-updates-to.html (As of 31 October 2018).

[221] Interview 11.

A select number of vendors affected by Spectre were initially notified on 1 June 2017, with subsequent details on Meltdown coming in July.[222] This initial group of vendors comprised major chip manufacturers Intel, AMD and ARM. The planned disclosure date for the two vulnerabilities was originally 7 January 2018, more than six months after the initial vulnerability report.

However, there were several hints of the vulnerabilities before Meltdown and Spectre were publicly disclosed, and critics argued that the multivendor coordination process broke down, as well as criticised the finders for leaving out key actors in the coordination process. Amazon, Google and Microsoft all issued patch notices before public disclosure, which hinted at the vulnerability, and a discussion on a Linux kernel email list prompted further discussion of the possible existence of a severe, yet undisclosed, vulnerability. Seven days before planned public disclosure, online news website The Register published a story on an 'Intel processor design flaw', ultimately leading the embargo to be lifted just a day later.[223] On 3 January 2018, the vulnerabilities were publicly disclosed, accompanied by a significant media presence, complete with a dedicated website, marketable abbreviations, logos, and associated Q&A and press material.[224]

The aftermath of the disclosure process also gained political attention, with US lawmakers questioning why certain vendors and organisations were left out of the disclosure process, as well as expressed geopolitical concerns on that Intel informed Chinese companies, including computer maker Lenovo, of the vulnerabilities before notifying CERT/CC or the US government.[225]

5.1.3 Meltdown and Spectre have had an extensive impact on manufacturers due to remediation costs rather than attack costs

As the vulnerability exists at the physical architecture stage, most important chip makers, including ARM, Intel and AMD, have been affected, meaning the vast majority of the world's laptops, desktops and smartphones are at risk from the two vulnerabilities.[226] Since the flaw was found to be so widely distributed, one of the finders, Daniel Grass, described it as, 'probably one of the worst CPU bugs ever found'.[227] Meltdown and Spectre also show the networked, interlinked nature of modern computing supply chains. However, the actually reported impacts due to exploitation have been limited and the economic consequences have so far mainly related to additional costs in developing, rolling out and implementing remediation measures. Nevertheless, this does not exclude the possibility of exploitation being uncovered in the future.

Further complications arose immediately after public disclosure and, since the disclosure timeline shifted to the left, many vendors were not fully prepared. Reports showed that some antivirus systems were caught off guard, inadvertently stopping the vendor patches from being deployed. Other patches rolled out had to be stopped mid-deployment after crashing machines.[228] Intel was perceived to handle the public disclosure poorly, initially downplaying the seriousness of the vulnerability and issuing a patch that was ill-received by the community. Linux creator Linus Torvalds even went as far as describing the first

[222] Gibbs (2018a).
[223] See Brandom (2018) for a more extensive discussion of the Meltdown and Spectre disclosure process.
[224] See https://meltdownattack.com/ (As of 31 October 2018).
[225] Hay Newman (2018).
[226] Coldewey (2018a).
[227] Gibbs (2018b).
[228] Brandom (2018).

patch as 'complete and utter garbage'.[229] After repeatedly committing to the quality of the initial patch, Intel then privately advised a select group to not to use the first patch without disclosing this to the public, which the company later did. As a result, Intel's shares fell nine per cent within the space of a week.[230] Additionally, Intel's CEO was accused of selling millions of dollars' worth of stock, between becoming aware of the vulnerabilities and public disclosure. Since disclosure, over 35 lawsuits have been initiated in relation to Intel's handling of the disclosure process.[231] However, Intel is not the only manufacturer that has faced criticism there are other companies as well that have suffered for both inadequate patches and reputational damages.

The consequences of the coordination process and difficulties in the initial patch deployment could have had significant downstream effects; businesses may incur higher costs since hardware will be less powerful, slower and need replacing sooner. Some reports claimed performance losses of as much as 30 per cent as a result of initial patching.[232] These costs are, however, difficult to quantify and, given the complexity of the vulnerability and the associated supply chain complexity, credit is due for vendors who were quick to deploy patches and mitigate any possible exploitation of the vulnerabilities.

5.1.4 The discovery of Meltdown and Spectre highlights the importance of investing in long-term security research

Beyond issues related to the disclosure process and its potential economic impacts, the Meltdown and Spectre also highlighted the importance and economic value of long-term and, oftentimes, fundamental security research. While CVD and bug bounty programmes may incentivise researchers to identify and report vulnerabilities in applications and services, they are less likely to identify complex and fundamental security vulnerabilities that require significant resources and technical skills to uncover.

All of the researchers who identified and ultimately disclosed Meltdown and Spectre received financial or organisational support to conduct their research. In the case of the academic researchers, their research was supported by their host institutions and European or American research grants. Jann Horn was enabled by Google Project Zero, whose mission is to seek to make the creation of software exploits more difficult and make zero-day exploits more costly.[233]

CVD thus plays an important part in the modern computing ecosystem, but it cannot be the only part; efforts must be made to develop more secure systems from start (i.e. 'secure by design') and to identify and mitigate current vulnerabilities in underlying infrastructure or enabling technologies that are unlikely to be mitigated by simply having a CVD policy or running a bug bounty programme.

[229] Coldewey (2018b).
[230] Financial Times (2018).
[231] https://www.businessinsider.com/intel-ceo-krzanich-sold-shares-after-company-was-informed-of-chip-flaw-2018-1
[232] The Wharton School of the University of Pennsylvania (2018).
[233] Interview 11.

5.2 EternalBlue

EternalBlue is a software vulnerability originally identified by the US National Security Agency (NSA), which was leaked to the public in 2017 by a hacker group known as 'The Shadow Brokers'. In contrast to Meltdown and Spectre, EternalBlue has been widely exploited and has incurred significant economic and societal costs.

EternalBlue is a vulnerability related to how the Windows Server Message Block (SMB) server handles particular requests. SMB is a networking protocol that is used for managing shared data between processes or providing shared access to services such as file storage, printers, and serial ports.[234] If successfully exploited, EternalBlue could enable a malicious actor to remotely execute arbitrary code on the target system.[235]

5.2.1 Discovery and disclosure of EternalBlue represents a non-disclosure process

In contrast to Meltdown and Spectre, EternalBlue was not subject to responsible disclosure. Rather, it was allegedly identified by the NSA and leaked to the public in April 2017.

The EternalBlue vulnerability was part of a larger NSA framework called 'FuzzBunch', which was designed to configure, deliver and execute exploits – much like the popular penetration testing framework 'Metasploit'.[236] EternalBlue, FuzzBunch and other custom-built NSA cyber capabilities were stolen from the NSA by The Shadow Brokers and subsequently leaked to the public in segments between August 2016 and April 2017.[237] The EternalBlue vulnerability and associated exploit were leaked by The Shadow Brokers as part of their fifth leak, which they labelled 'Lost in Translation' on 14 April 2017.

The initial release of the NSA-associated tools was announced by The Shadow Brokers in August 2016 in the form of an auction where the compromised data would be sold to the highest bidder through payment in the Bitcoin cryptocurrency.[238] The group emphasised that they were not seeking fame but rather a financial gain, and made several subsequent public posts in order to increase interest in purchasing the tools. The group further switched to direct sales due to limited engagement with the auction before announcing that they would stop posting publicly and delete their associated accounts due to the high risk involved.[239]

There is a suspicion that the NSA may have informed Microsoft of what vulnerabilities and tools 'The Shadow Brokers' stole once the leak was confirmed as legitimate. Microsoft issued an urgent security patch (MS17-010) that featured remediation measures for several of the SMB vulnerabilities exploited by EternalBlue and other leaked tools in March 2017, a full month before EternalBlue was leaked to the public.[240] NSA may have done so in order to limit the potential impact of the leaked vulnerability, but this has not been confirmed by the Agency.

[234] Microsoft (2018).
[235] Execution of arbitrary code would allow the malicious actor to run programs or take control over the target device, which is one of the most significant effects that can be achieved through the exploitation of a vulnerability.
[236] See https://www.metasploit.com/ (As of 31 October 2018).
[237] Sanger (2016).
[238] The Shadow Brokers (2016).
[239] The full post history of The Shadow Brokers can be viewed at https://steemit.com/@theshadowbrokers (As of 31 October 2018).
[240] Cimpanu (2018).

The initial disclosure of the NSA tools resulted in widespread concern within the US intelligence community since it was unclear who The Shadow Brokers were and how they had retrieved such top-secret material. The NSA also did not know the full scope of the breach, or which tools, techniques or vulnerabilities may have been compromised. 'The Shadow Brokers' also showed deep operational understanding about how the NSA worked, particularly in relation to the Tailored Access Operations (TAO) unit, one of its most elite units tasked with infiltrating high-profile targets.[241] This prompted concern at the presence of an insider within the NSA responsible for the leaks.

Reports in November 2017 made it clear that despite a wide-ranging, fifteen-month investigation by the NSA's counterintelligence unit and the Federal Bureau of Investigation (FBI), the NSA still does not seem to have confirmed how The Shadow Brokers gained access to their data. Three NSA employees have been arrested on suspicion of removing classified files since 2015, but none of them has been confirmed to be associated with The Shadow Brokers. There are allegations that The Shadow Brokers were directed and sponsored by Russian state institutions, but this has not been validated either.[242]

5.2.2 EternalBlue had significant and wide-ranging global implications

The Shadow Brokers disclosure has been perceived by many as being more harmful to US intelligence than the leak and publication of classified material by Edward Snowden in 2013.[243] First impressions of the leaked tools signalled that this was one of the most profound disclosures of vulnerabilities and exploits so far, particularly coming from a nation-state's intelligence service. However, many of the vulnerabilities were old and affected software and products that had since been upgraded or subject to new versions. This led to the view that organisations that followed information security good practice were unlikely to be affected, thereby considerably limiting the potential impact of the leaks.[244] Nevertheless, EternalBlue resulted in some of the most costly and destructive attacks and exploits in recent history. It also showcases the risk of moral hazard that may materialise in government disclosure decisions processes (GDDP), where decisions to not publicly disclose a vulnerability are made on national security grounds while subsequent costs of exploitation of that vulnerability are borne elsewhere.

Since the leak in April 2017, the EternalBlue vulnerability has been exploited by a number of different malware, including the fileless ransomware 'UIWIX', the SMB worm 'EternalRocks', the cryptocurrency mining malware 'Adylkuzz', and, most prominently, the 'WannaCry'[245] and 'Petya'/'NotPetya' ransomware.[246] In addition to EternalBlue, WannaCry also used the backdoor 'DoublePulsar' to propagate, which was another tool leaked by The Shadow Brokers. WannaCry infected over 300,000 computers in over 150 countries and is thought to have caused more than US$8 bn in damages.[247] It is widely believed that North Korea was behind the WannaCry attack, and Australia, Canada, New Zealand, the United Kingdom and the United States have all issued formal allegations against the North Korean government.[248]

The EternalBlue vulnerability was also used in the Petya ransomware and the associated destructive version NotPetya. While Petya performed similar functions as WannaCry, NotPetya only superficially

[241] Sanger (2016).
[242] Shane et al. (2017).
[243] Shane et al. (2017).
[244] Goodin (2017).
[245] WannaCry is a type of malware in the ransomware family that at the point of infection on a target machine encrypts user files, making them inaccessible, and demands a ransom payment in Bitcoin in order to decrypt them.
[246] Sanchez (2017), Sood & Hurley (2017).
[247] IBM X-Force (2018).
[248] BBC News (2017).

functioned as a ransomware and, instead, had much more destructive capabilities. Whereas a traditional ransomware allows for the decryption of files once a ransom has been paid, NotPetya also included functionality to permanently corrupt or destroy data.

The disclosure of EternalBlue has also led to further development of the vulnerability and associated exploits. There is now a publicly available version of EternalBlue that affects SMBv1 server deployments across a wide range of Windows versions, including Windows 7, 8 and 10.[249] This additional functionality allows for broader use and may enable new versions of malware for many different purposes, ranging from low-level crypto-mining operations to state-sponsored espionage or cyber attacks.[250] Despite the availability of security patches, there are still almost a million vulnerable Windows systems with exposed SMB services openly accessible through the Internet.[251]

5.2.3 The EternalBlue disclosure highlights the profound importance of responsible disclosure

The EternalBlue leak and subsequent attacks utilising the vulnerability have highlighted a number of important considerations. First, it highlights the important role that major technology vendors play in the event of significant security incidents. The speed at which major vendors develop and roll out appropriate remediation measures can have direct effects on limiting the impact and costs of attacks. Second, it also highlights the inability or lack of incentives for users to monitor security developments and apply appropriate security updates or patches – even when faced with significant threats. Many systems still remain vulnerable, even more than a year after initial patch release from Microsoft.

Third, it inadvertently reaffirms the importance of CVD and GDDP. The EternalBlue vulnerability enabled some of the most costly and destructive cyber attacks in recent history and could, in theory, have been avoided if it had been subjected to a CVD process or a more transparent GDDP process. The EternalBlue disclosure also highlights the challenging role that nation-states and governments play in the vulnerability disclosure landscape, and the significant consequences that non-disclosure can have if the vulnerability is leaked or stolen. The NSA claims that is has shared more than 90 per cent of vulnerabilities it has identified with the appropriate vendors, but this is difficult to verify without access to classified NSA data.[252] But in the light of EternalBlue, even the leak of any of the remaining 10 per cent may present such a significant threat of global harm that it may warrant further discussion of the ethics and responsibilities involved in the stockpiling of vulnerabilities, and the trade-offs involved with other social benefits (e.g. improved counterterrorism outcomes) pursued by intelligence agencies in doing so. Exactly what criteria a vulnerability must fulfil to be kept secret for national security purposes is still subject to ongoing debate. Nevertheless, the questions outlined in the 2017 US Government Vulnerabilities Equities Policy and Process (VEP), as shown in

[249] EternalBlue can now affect Microsoft Windows Vista SP2; Windows Server 2008 SP2 and R2 SP1; Windows 7 SP1; Windows 8.1; Windows Server 2012 Gold and R2; Windows RT 8.1; Windows 10 Gold, 1511, and 1607; and Windows Server 2016.
[250] Cimpanu (2018).
[251] According to a Shodan search for services running on port 445 with authentication disabled (June 2018).
[252] Menn (2018).

Table 5.1, may provide helpful guidance to what considerations could be made.

Table 5.1 Example equity considerations for GDDP from the US Government VEP

CONSIDERATION	EXAMPLE QUESTIONS
Part 1: Defensive equity considerations	
1.A. Threat considerations	• Where is the product used? How widely is it used?
1.B. Vulnerability considerations	• Is exploitation of this vulnerability alone sufficient to cause harm? How likely is it that threat actors will discover this vulnerability?
1.C. Impact considerations	• How severe is the vulnerability? What are the potential consequences of exploitation of this vulnerability?
1.D. Mitigation considerations	• If a patch is released, how likely is it to be applied to vulnerable systems? How soon? What percentage of vulnerable systems will remain unpatched for more than a year after patch release?
Part 2 – Intelligence, law enforcement and operational equity considerations	
2.A. Operational value considerations	• Can this vulnerability be exploited to support intelligence collection, cyber operations or law enforcement evidence collection?
2.B. Operational impact considerations	• Does exploitation of this vulnerability provide specialised operational value against cyber threat actors or their operations? Do alternative means exist to realize the operational benefits of exploiting this vulnerability?
Part 3 – Commercial equity considerations	• If United States Government (USG) knowledge of this vulnerability were to be revealed, what risks could that pose for USG relationships with industry?
Part 4 – International partnership equity considerations	• If USG knowledge of this vulnerability were to be revealed, what risks could that pose for USG international relations?

Source: United States Government (2017).

5.3 On reflection

The Meltdown/Spectre and EternalBlue vulnerability disclosures illustrate the extent of modern society's dependence on computers and connectivity. The two case studies also highlight the fragility of this ecosystem, where widely used applications, systems and enabling technologies can contain significant and exploitable vulnerabilities capable of causing significant economic and other societal harm. Both case studies further underline the interconnected nature of the global ICT supply chain in which vulnerabilities are found both horizontally and vertically – increasing the demands on coordinators' and vendors' abilities to perform multivendor CVD. The Meltdown and Spectre case study also highlights the liability issues present in the modern computing market. While software updates and patches may provide partial protection from the vulnerabilities, full protection is likely to require the replacement of vulnerable hardware. The wide-spread replacement of CPUs will incur significant costs to users, which is unlikely to be borne by the chip manufacturers.[253]

While Meltdown and Spectre certainly presented a unique challenge to the chip-making industry and its dependencies, it is cannot be discarded that other vulnerabilities with possible wide-ranging implications can be discovered in other segments of the computing industry in the near future. Considering the growth in the number of deployed CPUs and the growth of Internet-enabled IoT devices and other cyber-physical systems, it is most likely simply a matter of time before another vulnerability with systemic impact potential materialises.

As seen in both the case studies, the vulnerability disclosure ecosystem is still evolving and maturing, particularly in relation to complex multivendor CVD and GDDPs. Many of the key aspects of multivendor CVD remain open, including how to decide which vendors to involve in the process and when, how far down the supply chain the coordination process should extend to, as well as if and which government agencies should be involved in the process.

Lastly, it is clear that the potential impacts that vulnerabilities can impose on society, whether responsibly disclosed or not, shows that the economics of vulnerabilities extend far beyond just vulnerability disclosure. The examination of the societal costs for information and network security must therefore also be approached comprehensively and holistically in order to incentivise an environment that promotes the secure development of software, hardware and services that also is capable of identifying, reporting and mitigating vulnerabilities that materialise.

[253] Hay Newman (2018b).

6. Summary and key findings

6.1 Summary

This study set out to explore the economic parameters and incentives across the different actors within the vulnerability disclosure lifecycle in order to help improve vulnerability disclosure processes and programmes. Chapter 2 presented the different actors and processes within vulnerability disclosure and discussed how they may engage with each other through the different vulnerability disclosure processes. Chapter 3 further illustrated that vulnerability disclosure takes place in a wider computing and information security ecosystem whose unique economic structures and incentives have direct economic effects on vulnerability disclosure. This extends to both the structure and parameters of the market for information security as well as the nature of software and hardware vulnerabilities. This wider ecosystem is also subject to continuous change driven by technological development to which vulnerability disclosure must adapt and respond to. In the coming years, there will be a continued increase in deployment of cyber-physical and IoT systems of relatively poor security where exploitable vulnerabilities could result in significant societal harm.

Chapter 4 examined the behaviour and incentives at the individual, organisational and structural level in the vulnerability disclosure ecosystem in further detail. It is clear that economic incentives play a key role in vulnerability disclosure across all actors and processes, regardless of what type of vulnerability disclosure process is ultimately pursued. This emphasises the importance of a well-developed understanding of what the economic parameters in vulnerability disclosure are and how they may influence different processes. Chapter 4 also highlighted that many of the behaviours and incentives in vulnerability disclosure are affected by both negative and positive external factors, emphasising that behaviour in vulnerability disclosure can be influenced and changed through different mechanisms.

Chapter 5 featured two case studies of recently disclosed high-profile vulnerabilities and illustrated how vulnerability disclosure may take place in practice and how economic parameters or incentives may be realised. The two case studies highlighted the diversity found in vulnerability disclosure and illustrated the distinct, and potentially grave, differences between a CVD and a non-disclosure process. While both case studies showed the costs that vulnerabilities can incur, they also illustrated the cost savings that responsible disclosure can realise by reducing the exploitation of identified vulnerabilities. The Meltdown and Spectre case study also showed the significant resources that are needed to conduct a CVD process for complex vulnerabilities with wide-ranging supply chain consequences, as well as the significant effort needed to develop and roll out patches for such vulnerabilities. Lastly, both case studies emphasised the liability challenges present in vulnerability disclosure and information security, where there is still uncertainty as to how and to what extent the vulnerability identifier (NSA in the case of EternalBlue) or vulnerability owner (the chip makers in the case of Meltdown and Spectre) can be held accountable for costs incurred by other parties.

6.2 Key findings

Overall, the study has a produced a number of key findings. First and foremost, the study shows the importance that vulnerability disclosure, and predominantly CVD, plays in modern society. As witnessed in the case of EternalBlue, vulnerabilities in widely used software and hardware can cause immense societal harm across the globe and it is necessary to have processes in place to adequately identify, report, receive, triage and mitigate vulnerabilities. As the potential risk is so severe, vendors that develop or manufacture products or services for the Internet or the global ICT ecosystem may no longer have the choice to not have the ability to receive good-faith vulnerability reports from the community. In the modern computing ecosystem, everybody will have vulnerabilities, so being able to receive them and respond to them is what ultimately matters most. Similarly, national governments should adopt a CVD policy and begin a discussion of how to best approach a government disclosure decisions process.

The study findings also emphasises the importance of approaching vulnerability disclosure as an ecosystem. CVD requires a finder and a vendor (and sometimes a coordinator), and the success of a CVD process rests on the relationships between these actors. All actors involved in vulnerability disclosure should therefore recognise the importance of setting up and running mutually beneficial structures that enables effective and efficient CVD to take place. Awareness raising and capacity building across all actor groups are key enablers for this to happen and for actors to understand the economic incentives and behaviour of other parties involved in CVD. Providing actors with resources, good practice and voluntary standards are also important tools to consider in promoting mutually beneficial and standardised behaviour. Communication skills are also critically important in CVD; finders and vendors alike must be able to constructively engage with each other in a timely fashion and in a shared language that both parties understand. This type of ecosystem thinking also extends to how finders are treated by vendors and coordinators. Most prominently, there are opportunities to improve finder wellbeing and the overall CVD ecosystem by ensuring safe harbour practices and legal safeguards for security researchers working to identify and report vulnerabilities.

However, the study has also reaffirmed that the economics of vulnerability disclosure is an emerging area of research. Much of the evidence cited in this report is anecdotal or theoretical and there is a clear need for additional empirical data and statistically relevant, longitudinal research in a number of areas, including:

- **The motivations of finders**, particularly extending beyond bug bounty programmes. CVD programmes are a fundamental component of modern information security and it is important to have a clear understanding of why certain security researchers work to identify and report vulnerabilities, even in the absence of financial compensation.
- **How to better quantify the cost of the exploitation of vulnerabilities** to inform discussions on liability, insurance and other structural levers (to inform liability, insurance, etc.). It is currently challenging to show the correlation between poor security practices or non-existent vulnerability disclosure processes and societal harm caused by the exploitation of vulnerabilities.
- **The cost of implementing and running vulnerability disclosure programmes** to help organisations make better informed decisions about security investments and trade-offs between different types of security interventions. Some organisations may be better suited to running a particular type of vulnerability disclosure programme, whereas other organisations are not yet ready to implement any vulnerability disclosure intervention.
- **Quantification of security gains through vulnerability disclosure** to better understand the value that vulnerability disclosure programme bring (i.e. are they a worthwhile security investment?). This could also help in building businesses cases for CVD adoption that could be used to show less mature organisations the value of having a CVD policy or programme.
- **The cost of developing and implementing patches** to understand the economic costs associated with vulnerabilities beyond costs accrued due to their exploitation. This would entail an examination of the complexity, time and resources required to develop patches for different types of vulnerabilities, as well as an exploration what costs organisations incur in the roll out and implementation of different types of patches for different types of software and hardware.

A well-developed understanding of all of the incentives and economic parameters mentioned above is crucial to develop appropriate CVD programmes and in order to not over- or under-incentivise any segments of vulnerability disclosure market.

Finally, vendors and other organisations should be mindful of the difference between CVD, bug bounties and other information security interventions. Most organisations should consider implementing a CVD process, and some may want to consider a bug bounty programme, but not at the cost of other information security interventions in the development and testing stage (e.g. code audits, penetration testing, vulnerability assessments, etc.). CVD and bug bounty programmes also typically only identify possible vulnerabilities in systems and services once they are live and operational. Continuous efforts should also be made in improving the quality and security of software and hardware throughout the development lifecycle so as to reduce the number of vulnerabilities in deployment.

Lastly, while CVD and bug bounty programmes may be able to identify certain types of vulnerabilities, they are unlikely to identify larger structural issues present in modern computing systems. Governments, academic instructions and private organisations should therefore keep investing in long-term security research in both academic and private sector settings in order to identify and mitigate fundamental weaknesses such as design flaws or protocol vulnerabilities.

7. References

Ablon, Lillian, Martin C. Libicki, and Andrea A. Golay. 2014. 'Markets for Cybercrime Tools and Stolen Data: Hackers' Bazaar.' Santa Monica, Calif.: RAND Corporation, RR-610-JNI As of October 25, 2018: https://www.rand.org/pubs/research_reports/RR610.html

Ablon, Lillian and Andy Bogart. 2017. 'Zero Days, Thousands of Nights: The Life and Times of Zero-Day Vulnerabilities and Their Exploits.' Santa Monica, Calif.: RAND Corporation, RR-1751-RC. As of October 25, 2018: https://www.rand.org/pubs/research_reports/RR1751.html

Acquisti, Alessandro, & Jens Grossklags. 2005. 'Privacy and rationality in individual decision making.' *IEEE Security & Privacy* 3(1): 26–33.

Akerloff, George. 1970. 'The market for lemons: Quality uncertainty and the market mechanism.' *Quarterly journal of economics* 84(3): 488-500.

Algarni, Abdullah M, & Yashwant K Malaiya. 2013. 'Most successful vulnerability discoverers: Motivation and methods.' Proceedings of the International Conference on Security and Management (SAM). The Steering Committee of The World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp).

———. 2014. 'Software vulnerability markets: Discoverers and buyers.' *International Journal of Computer, Information Science and Engineering* 8(3): 71–81.

Anderson, Ross. 2001. 'Why information security is hard-an economic perspective.' In Computer security applications conference, 2001. Acsac. proceedings 17th annual, pp. 358_365. *IEEE*, 2001.

Anderson, Ross, & Tyler Moore. 2006. 'The economics of information security.' *Science* 314(5799): 610_3.

Anderson, Ross, Rainer Böhme, Richard Clayton & Tyler Moore. 2009. 'Security economics and European policy.' In *Managing information risk and the economics of security*, 55-80. Springer, Anderson, Ross. 2002. 'Security in open versus closed systems—the dance of Boltzmann, Coase and Moore.' Technical report, Cambridge University, England.

Arora, Ashish, & Rahul Telang. 2005. 'Economics of software vulnerability disclosure.' *IEEE Security & Privacy* 3(1): 20–5.

Arora, Ashish, Rahul Telang & Hao Xu. 2008. 'Optimal policy for software vulnerability disclosure.' *Management Science* 54(4): 642–56.

Arora, Ashish, Ramayya Krishnan, Rahul Telang & Yubao Yang. 2010. 'An empirical analysis of software vendors' patch release behavior: impact of vulnerability disclosure.' *Information Systems Research* 21(1): 115–32.

Baer, Walter S, & Andrew Parkinson. 2007. 'Cyberinsurance in IT security management.' *IEEE Security & Privacy* 5(3).

Bailey, Russell, & Barbara Tierney. 2002. 'Information commons redux: concept, evolution, and transcending the tragedy of the commons.' *The Journal of Academic Librarianship* 28(5): 277–86.

BBC News. 2017. 'Cyber-attack: US and UK blame North Korea for WannaCry.' As of 23 June 2018: https://www.bbc.com/news/world-us-canada-42407488

Böhme, R. 2005. 'Vulnerability markets—What is the economic value of a zero-day exploit?' Paper held at the 2005 Chaos Communication Congress Berlin, Germany.

Böhme, Rainer, & Gaurav Kataria. 2006. 'Models and Measures for Correlation in Cyber-Insurance.' WEIS.

Böhme, Rainer, & Galina Schwartz. 2010. 'Modeling Cyber-Insurance: Towards a Unifying Framework.' WEIS.

Böhme, Rainer. 2010. 'Towards insurable network architectures.' *it-Information Technology Methoden und innovative Anwendungen der Informatik und Informationstechnik* 52(5): 290-3.

Brandom, Russell. 2018. 'Keeping Spectre Secret: How an industry-breaking bug stayed secret for seven months — and then leaked out.' As of 31 October 2018: https://www.theverge.com/2018/1/11/16878670/meltdown-spectre-disclosure-embargo-google-microsoft-linux

Bright, Peter. 2018. 'Meltdown and Spectre: Here's what Intel, Apple, Microsoft, others are doing about it.' As of 31 October 2018: https://arstechnica.com/gadgets/2018/01/meltdown-and-spectre-heres-what-intel-apple-microsoft-others-are-doing-about-it/

Bugcrowd. n.d. 'Defensive Vulnerability Pricing Model: How to budget for your crowdsourced security program.' As of 31 October 2018: https://cdn2.hubspot.net/hubfs/1549768/PDFs/Whats_a_Bug_Worth.pdf

———. 2018. '2018 State of Bug Bounty.' As of 25 October 2018: https://www.bugcrowd.com/resource/2018-state-of-bug-bounty-report/

Camp, L Jean, & Catherine Wolfram. 2004. 'Pricing security.' In *Economics of information security*, 17–34. Springer.

Cavusoglu, Hasan, Huseyin Cavusoglu & Srinivasan Raghunathan. 2005. 'Emerging Issues in Responsible Vulnerability Disclosure.' WEIS.

Centre for European Policy Studies (CEPS). 'Software Vulnerability Disclosure in Europe: Technology, Policies and Legal Challenges. Report of a CEPS Task Force.' CEPS Task Force Reports 28 June 2018. As of 31 October 2018: https://www.ceps.eu/content/software-vulnerability-disclosure-europe

CERT/CC. 2017a. 'Vulnerability Disclosure Policy.' As of 31 October 2018: http://www.cert.org/vulnerability-analysis/vul-disclosure.cfm?

———. 2017b. 'Vulnerability Note VU#307015.' As of 31 October 2018: https://www.kb.cert.org/vuls/id/307015

———. 2017c. 'Vendor Information for VU#228519.' As of 31 October 2018: https://www.kb.cert.org/vuls/byvendor?searchview&Query=FIELD+Reference=228519&SearchOrder=4

Christin, Nicolas, Serge Egelman, Timothy Vidas & Jens Grossklags. 2011. 'It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice.' International Conference on Financial Cryptography and Data Security. Springer.

Cimpanu, Catalin. 2018. ‘One Year After WannaCry: EternalBlue Exploit is Bigger Than Ever.’ As of 31 October 2018: https://www.bleepingcomputer.com/news/security/one-year-after-wannacry-eternalblue-exploit-is-bigger-than-ever/

Coldewey, Devin. 2018a. ‘Kernel panic! What are Meltdown and Spectre, the bugs affecting nearly every computer and device?’ As of 31 October 2018: https://techcrunch.com/2018/01/03/kernel-panic-what-are-meltdown-and-spectre-the-bugs-affecting-nearly-every-computer-and-device/

———. 2018b. ‘Linus Torvalds declares Intel fix for Meltdown/Spectre “COMPLETE AND UTTER GARBAGE”.’ As of 31 October 2018: https://techcrunch.com/2018/01/22/linus-torvalds-declares-intel-fix-for-meltdown-spectre-complete-and-utter-garbage/

Day, Oliver, Brandon Palmen & Rachel Greenstadt. 2009. ‘Reinterpreting the disclosure debate for web infections.’ In *Managing information Risk and the Economics of Security*, 179–97. Springer.

Delcheva, Teodora & Soesanto, Stefan. 2018. ‘Time to talk: Europe and the Vulnerability Equities Process.’ As of 31 October: https://www.ecfr.eu/article/commentary_time_to_talk_europe_and_the_vulnerability_equities_process

Dullien, Thomas. 2018. ‘Security, Moore’s law and the anomaly of cheap complexity. CYCON.’ As of 31 October 2018: https://www.err.ee/836236/video-google-0-projekti-tarkvarainseneri-ettekanne-cyconil

Edmundson, Anne, Brian Holtkamp, Emanuel Rivera, Matthew Finifter, Adrian Mettler & David Wagner. 2013. ‘An empirical study on the effectiveness of security code review.’ International Symposium on Engineering Secure Software and Systems. Springer.

Elazari Bar On, Amit. 2018. ‘Private Ordering Shaping Cybersecurity Policy – The Case of Bug Bounties’, forthcoming in *Rewired: The Past, Present, and Future of Cybersecurity* (Ryan Ellis & Vivek Mohan Eds., 2018).

ENISA. 2015. ‘Good Practice Guide on Vulnerability Disclosure: From challenges to recommendations.’ As of 31 October 2018: https://www.enisa.europa.eu/publications/vulnerability-disclosure

———. 2018. ‘ENISA Regulatory Framework.’ As of 31 October 2018: https://www.enisa.europa.eu/about-enisa/regulatory-framework

European Council. 2017. Council Conclusions on the Joint Communication to the European Parliament and the Council, Resilience, Deterrence and Defence: Building strong cybersecurity for the EU – Council conclusions (20 November 2017), Conclusion No. 27.

Financial Times (2018). ‘Markets Data – Equities.’ As of 31 October 2018: https://markets.ft.com/data/equities/tearsheet/summary?s=INTC:NSQ

Finifter, Matthew, Devdatta Akhawe & David Wagner. 2013. *An Empirical Study of Vulnerability Rewards Programs*. USENIX Security Symposium.

FIRST. 2017. ‘Multiparty Vulnerability Disclosure.’ As of 31 October 2018: https://www.first.org/global/sigs/vulnerability-coordination/multiparty/FIRST-Multiparty-Vulnerability-Coordination-v1.0.pdf

Frei, Stefan, & Francisco Artes. 2013. ‘International Vulnerability Purchase Program: Why buying all vulnerabilities above black market prices is economically sound.’ NSS Labs, December.

Gibbs, Samuel. 2018a. ‘Meltdown and Spectre: “worst ever” CPU bugs affect virtually all computers’ (2018). As of 31 October 2018: https://www.theguardian.com/technology/2018/jan/04/meltdown-spectre-worst-cpu-bugs-ever-found-affect-computers-intel-processors-security-flaw

———. 2018b. ‘Spectre and Meltdown processor security flaws – explained.’ As of 31 October 2018: https://www.theguardian.com/technology/2018/jan/04/meltdown-spectre-computer-processor-intel-security-flaws-explainer

Global Information Security Workforce Study (GISWS). 2017. ‘2017 Global Information Security Workforce Study.’ As of 31 October 2018: https://iamcybersafe.org/wp-content/uploads/2017/06/Europe-GISWS-Report.pdf

Goodin, Dan. 2017. ‘NSA-leaking Shadow Brokers just dumped its most damaging release yet.’ As of 31 October 2018: https://arstechnica.com/information-technology/2017/04/nsa-leaking-shadow-brokers-just-dumped-its-most-damaging-release-yet/

Gordon, Lawrence A., and Martin P. Loeb. ‘The economics of information security investment.’ *ACM Transactions on Information and System Security* (TISSEC) 5, no. 4 (2002): 438-457.

Granick, Jennifer S. 2005. ‘Faking It: Criminal Sanctions and the Cost of Computer Intrusions.’ WEIS.

Greenberg, Andy. 2018. ‘Triple Meltdown: How So Many Researchers Found a 20-Year-Old Chip Flaw At the Same Time.’ As of 31 October 2018: https://www.wired.com/story/meltdown-spectre-bug-collision-intel-chip-flaw-discovery/

Grossklags, Jens, Nicolas Christin & John Chuang. 2008. ‘Secure or insure?: a game-theoretic analysis of information security games.’ Proceedings of the 17th international conference on World Wide Web. ACM.

HackerOne. 2017. ‘The Hacker-powered Security Report 2017.’ As of 31 October 2018: https://www.hackerone.com/sites/default/files/2017-06/The%20Hacker-Powered%20Security%20Report.pdf

———. 2017b. ‘Key Reinstallation Attacks: Breaking WPA2 by forcing nonce reuse.’ As of 31 October 2018: https://hackerone.com/reports/286740

———. 2018. ‘The Hacker-powered Security Report 2018.’ As of 31 October 2018: https://www.hackerone.com/resources/hacker-powered-security-report

Hahn, Robert W, & Anne Layne-Farrar. 2006. ‘The law and economics of software security.’ AEI-Brookings Joint Center Working Paper No. 06-08. Available at SSRN: https://ssrn.com/abstract=897725 or http://dx.doi.org/10.2139/ssrn.897725

Hardin, Garrett. 2009. ‘The tragedy of the commons.’ *Journal of Natural Resources Policy Research* 1(3): 243–53.

Hay Newman, Lily. 2018a. ‘Senators Fear Meltdown and Spectre Disclosure Gave China An Edge.’ As of 31 October 2018: https://www.wired.com/story/meltdown-and-spectre-intel-china-disclosure/

———. 2018b. ‘After Meltdown and Spectre, another scary chip flaw emerges.’ As of 31 October: https://www.wired.com/story/speculative-store-bypass-spectre-meltdown-vulnerability/

He, Changhua, Mukund Sundararajan, Anupam Datta, Ante Derek & John C Mitchell. 2005. 'A modular correctness proof of IEEE 802.11 i and TLS.' Proceedings of the 12th ACM conference on Computer and communications security. ACM.

Householder, Allen D, Garret Wassermann, Art Manion & Chris King. 2017. 'The CERT® Guide to Coordinated Vulnerability Disclosure.' As of 31 October 2018: https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=503330

I am the Cavalry. n.d. '5 Motivations of Security Researchers.' As of 31 October 2018: https://www.iamthecavalry.org/motivations

IBM X-Force. 2018. 'IBM X-Force Threat Intelligence Index.' As of 31 October 2018: https://securityintelligence.com/2018-ibm-x-force-report-shellshock-fades-gozi-rises-and-insider-threats-soar/

Jaquith, Andrew. 2007. *Security metrics: replacing fear, uncertainty, and doubt*. Pearson Education.

Johnson, Benjamin, Rainer Böhme & Jens Grossklags. 2011. 'Security games with market insurance.' International Conference on Decision and Game Theory for Security. Springer.

Kannan, Karthik, & Rahul Telang. 2005. 'Market for software vulnerabilities? Think again.' Management Science 51(5): 726-40.

Kesan, Jay, Ruperto Majuca & William Yurcik. 2005. 'Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study.' Proc. WEIS.

Kocher, Paul, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz & Yuval Yarom. 2018. 'Spectre Attacks: Exploiting Speculative Execution.' arXiv preprint arXiv:1801.01203.

Krugman, P. 2009. *The Return of Depression Economics and the Crisis of 2008*. W. W. Norton.

Kunreuther, Howard, & Geoffrey Heal. 2003. 'Interdependent security.' *Journal of risk and uncertainty* 26(2–3): 231–49.

Lakhani, Aamir. 2017. 'How does WannaCry spread?' As of 31 October 2018: https://www.fortinet.com/blog/threat-research/wannacry-faq.html

Laszka, Aron, Mingyi Zhao & Jens Grossklags. 2016. 'Banishing misaligned incentives for validating reports in bug-bounty platforms.' European Symposium on Research in Computer Security. Springer.

Leverett, Éireann, Richard Clayton, and Ross Anderson. 2017. 'Standardisation and Certification of the "Internet of Things".' Proceedings of WEIS 2017 (2017). As of 31 October 2018: https://www.cl.cam.ac.uk/~rja14/Papers/eu-jrc-77862.pdf

Lewis, Paul Simon. 2017. 'The global vulnerability discovery and disclosure system: a thematic system dynamics approach.' As of 31 October 2018: http://dspace.lib.cranfield.ac.uk/handle/1826/12665

Li, Pu, & H Raghav Rao. 2007. 'An examination of private intermediaries' roles in software vulnerabilities disclosure.' *Information Systems Frontiers* 9(5): 531–9.

Lipp, Moritz, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom & Mike Hamburg. 2018. 'Meltdown.' arXiv preprint arXiv:1801.01207.

Maillart, Thomas, Mingyi Zhao, Jens Grossklags & John Chuang. 2017. 'Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs.' *Journal of Cybersecurity* 3(2): 81–90.

McKinney, David. 2007. 'Vulnerability bazaar.' *IEEE Security & Privacy* 5(6).

Meltdown and Spectre. 2018. 'Meltdown and Spectre'. As of 31 October 2018: https://meltdownattack.com/

Menn, Joseph. 2018. 'NSA says how often, not when, it discloses software flaws.' As of 31 October 2018: https://www.reuters.com/article/us-cybersecurity-nsa-flaws-insight/nsa-says-how-often-not-when-it-discloses-software-flaws-idUSKCN0SV2XQ20151107

Microsoft. 2018. 'Microsoft SMB Protocol and CIFS Protocol Overview.' As of 31 October 2018: https://docs.microsoft.com/en-us/windows/desktop/FileIO/microsoft-smb-protocol-and-cifs-protocol-overview

Miller, Charlie. 2007. 'The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales.' In Sixth Workshop on the Economics of Information Security. Citeseer.

Moore, Tyler, Allan Friedman & Ariel D Procaccia. 2010. 'Would a "cyber warrior" protect us: exploring trade-offs between attack and defense of information systems.' *Proceedings of the 2010 New Security Paradigms Workshop. ACM.*

Mullin, Wallace P. 2001. 'Will gun buyback programs increase the quantity of guns?' *International Review of Law and Economics* 21(1): 87–102.

Munaiah, Nuthan, & Andrew Meneely. 2016. 'Vulnerability severity scoring and bounties: why the disconnect?' Proceedings of the 2nd International Workshop on Software Analytics. ACM.

Nakashima, Ellen, & Craig Timberg. 2017. 'NSA officials worried about the day its potent hacking tool would get loose. Then it did.' Washington Post. As of 31 October 2018: https://www.washingtonpost.com/business/technology/nsa-officials-worried-about-the-day-its-potent-hacking-tool-would-get-loosethen-it-did/2017/05/16/50670b16-3978-11e7-a058-ddbb23c75d82_story.html

National Telecommunications and Information Administration (NTIA). 2016. 'Vulnerability Disclosure Attitudes and Actions: A Research Report from the NTIA Awareness and Adoption Group.' As of 31 October 2018: https://www.ntia.doc.gov/files/ntia/publications/2016_ntia_a_a_vulnerability_disclosure_insights_report.pdf

Neuhaus, Stephan, and Bernhard Plattner. 'Software security economics: Theory, in practice.' *The Economics of Information Security and Privacy*, pp. 75-92. Springer, Berlin, Heidelberg.

Ozment, James Andrew. 2007. *Vulnerability discovery & software security*. University of Cambridge.

Peeters, Gijs. 2017. 'Strengthening the digital Achilles heel of the European Union: Make use of ethical hackers to find vulnerabilities in information systems?' As of 31 October 2018: https://openaccess.leidenuniv.nl/handle/1887/55426

Radianti, Jaziar, & Jose J Gonzalez. 2007. 'A preliminary model of the vulnerability black market.' 25th International System Dynamics Conference at Boston, USA.

Radianti, Jaziar. 2010. 'Eliciting information on the vulnerability black market from interviews.' Emerging Security Information Systems and Technologies (SECURWARE), 2010 Fourth International Conference on. IEEE.

Rahman, Akond Ashfaque Ur, & Laurie Williams. 2016. 'Software security in devops: synthesizing practitioners' perceptions and practices.' Continuous Software Evolution and Delivery (CSED), IEEE/ACM International Workshop on. IEEE.

Ransbotham, Sam, Sabyaschi Mitra & Jon Ramsey. 2012. 'Are markets for vulnerabilities effective?' *Mis Quarterly*: 43–64.

Rashid, Fahmida. 2018. 'Web Application Bugs, from Disclosure to Exploit'. As of 31 October 2018: https://duo.com/decipher/web-application-bugs-from-disclosure-to-exploit

Rescorla, Eric. 2005. 'Is finding security holes a good idea?' *IEEE Security & Privacy* 3(1): 14-9.

Romanosky, Sasha, Lillian Ablon, Andreas Kuehn, and Therese Jones. 2017. 'Content Analysis of Cyber Insurance Policies: How do carriers write policies and price cyber risk?'. Santa Monica, Calif.: RAND Corporation, WR-1208, 2017. As of 31 October 2018: https://www.rand.org/pubs/working_papers/WR1208.html

Ruohonen, Jukka, & Luca Allodi. 2018. 'A Bug Bounty Perspective on the Disclosure of Web Vulnerabilities.' arXiv preprint arXiv:1805.09850.

Sanchez, William Gamazo. 2017. 'MS17-010: EternalBlue's Large Non-Paged Pool Overflow in SRV Driver.' As of 31 October 2018: https://blog.trendmicro.com/trendlabs-security-intelligence/ms17-010-eternalblue/

Sanger, David. 2016. '"Shadow Brokers" Leak Raises Alarming Question: Was the N.S.A. Hacked?' As of 31 October 2018: https://www.nytimes.com/2016/08/17/us/shadow-brokers-leak-raises-alarming-question-was-the-nsa-hacked.html

Shahzad, Muhammad, Muhammad Zubair Shafiq & Alex X Liu. 2012. 'A large scale exploratory analysis of software vulnerability life cycles.' Proceedings of the 34th International Conference on Software Engineering. IEEE Press.

Shane, Scott, Nicole Perlroth and David E. Sanger. 2017. 'Security Breach and Spilled Secrets Have Shaken the N.S.A. to Its Core.' As of 31 October 2018: https://www.nytimes.com/2017/11/12/us/nsa-shadow-brokers.html

Sood, Karan & Shaun Hurley. 2017. 'NotPetya Technical Analysis – A Triple Threat: File Encryption, MFT Encryption, Credential Theft.' As of 31 October 2018: https://www.crowdstrike.com/blog/petrwrap-ransomware-technical-analysis-triple-threat-file-encryption-mft-encryption-credential-theft/

Takanen, Ari, Petri Vuorijärvi, Marko Laakso & Juha Röning. 2004. 'Agents of responsibility in software vulnerability processes.' *Ethics and Information Technology* 6(2): 93–110.

Tang, MingJian, Mamoun Alazab & Yuxiu Luo. 2016. 'Exploiting vulnerability disclosures: statistical framework and case study.' Cybersecurity and Cyberforensics Conference (CCC), 2016. IEEE.

Temizkan, Orcun, Ram L Kumar, Sungjune Park & Chandrasekar Subramaniam. 2012. ‘Patch release behaviors of software vendors in response to vulnerabilities: an empirical analysis.’ *Journal of management information systems* 28(4): 305–38.

The Netherlands National Cybersecurity Centre. 2018. ‘Coordinated Vulnerability Disclosure: The Guideline.’ As of 31 October 2018: https://www.ncsc.nl/english/current-topics/responsible-disclosure-guideline.html

The Shadow Brokers. 2016. As of 31 October 2018: https://steemit.com/shadowbrokers/@theshadowbrokers/repost-theshadowbrokers-message-1-august-2016

The Wharton School of the University of Pennsylvania. 2018. ‘How Spectre and Meltdown Will Impact Companies and Consumers.' As of 31 October 2018: http://knowledge.wharton.upenn.edu/article/spectre-and-meltdown/

United States Government. 2017. ‘Vulnerabilities Equities Policy and Process for the United States Government’. As of 31 October 2018: https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

Vaas, Lisa. 2018. ‘“Misguided” hacking bill threatens to ice security researchers, say critics.’ As of 31 October 2018: https://nakedsecurity.sophos.com/2018/02/28/misguided-hacking-bill-threatens-to-ice-security-researchers-say-critics/

Van Eeten, Michel J, & Johannes M Bauer. 2008. ‘Economics of malware: Security decisions, incentives and externalities.’ OECD Science, Technology and Industry Working Papers 2008(1): 0_1.

Verizon. 2018. ‘Verizon 2018 Data Breach Investigations Report.’ As of 31 October 2018: http://www.verizonenterprise.com/verizon-insights-lab/dbir/

Wolverton, Troy. 2018. ‘Amid controversy over Intel CEO's stock sale, SEC warns executives about trading shares before disclosing security breaches.’ As of 31 October 2018: http://www.businessinsider.com/sec-issue-guidelines-regarding-disclosure-of-security-breaches-2018-2?r=UK&IR=T

Zero-day Initiative. n.d. ‘Disclosure Policy.’ As of 31 October 2018: http://www.zerodayinitiative.com/advisories/disclosure_policy/

Zhao, Mingyi, Jens Grossklags & Kai Chen. 2014. ’An exploratory study of white hat behaviors in a web vulnerability disclosure program.’ Proceedings of the 2014 ACM workshop on security information workers. ACM.

Zhao, Mingyi, Jens Grossklags & Peng Liu. 2015. ‘An empirical study of web vulnerability discovery ecosystems.’ Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. ACM.

Zhao, Mingyi, Aron Laszka, Thomas Maillart & Jens Grossklags. 2016. ‘Modeling and organizing bug-bounty programs.’ The HCOMP Workshop on Mathematical Foundations of Human Computation, Austin, TX, USA.

Zhao, Mingyi, Aron Laszka & Jens Grossklags. 2017. ‘Devising effective policies for bug-bounty platforms and security vulnerability discovery.’ *Journal of Information Policy* 7: 372 –418.

Annex A: List of interviewees

#	NAME	TITLE AND ORGANISATION
1	Joern Schneeweisz	Vulnerability Researcher, Recurity Labs
2	Casey Ellis	Chairman, Founder & CTO, Bugcrowd
3	Rainer Boehme	Professor for Security and Privacy, Department of Computer Science, University of Innsbruck, Austria
4	Dirk Jumpertz	Security Manager, EURid
5	Peter Allor	Sr. Director, Chief Security Officer & Product Cyber Security Chief Honeywell Connected Enterprise
6	Shamiq Islam	Application Security Manage, Coinbase
7	Stijn Jans	Founder, Intigriti
8	Alex Rice	CTO, HackerOne
9	Jeroen van der Ham	Assistant professor/researcher, University of Twente, The Netherlands National Cyber Security Centre
10	James Ritchey	Senior Application Security Engineer, Gitlab
11	Ben Hawkes	Google Project Zero
12	Allen Householder	Senior Vulnerability Analyst, CERT/CC at Carnegie Mellon University's Software Engineering Institute
13	Amit Elazari	UC Berkeley School of Law, UC Berkeley Center for Long-Term Cybersecurity, and UC Berkeley School of Information's Master of Information and Cybersecurity

ENISA

European Union Agency for Network
and Information Security
1 Vasilissis Sofias
Marousi 151 24, Attiki, Greece

Heraklion Office

Science and Technology Park of Crete (ITE)
Vassilika Vouton, 700 13, Heraklion, Greece

Catalogue Number: TP-07-18-080-EN-N

1 Vasilissis Sofias Str, Maroussi 151 24, Attiki, Greece
Tel: +30 28 14 40 9710
info@enisa.europa.eu
www.enisa.europa.eu

ISBN: 978-92-9204-277-6
DOI: 610.2824/49807