



Delft University of Technology

One thing after another

The role of users, manufacturers, and intermediaries in iot security

Turcios Rodriguez, E.R.

DOI

[10.4233/uuid:64e15692-06d7-4e3a-9d51-97f4a07b403f](https://doi.org/10.4233/uuid:64e15692-06d7-4e3a-9d51-97f4a07b403f)

Publication date

2023

Document Version

Final published version

Citation (APA)

Turcios Rodriguez, E. R. (2023). *One thing after another: The role of users, manufacturers, and intermediaries in iot security*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:64e15692-06d7-4e3a-9d51-97f4a07b403f>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

ONE
(T)

H

I

N

G

AFTER ANOTHER

THE ROLE OF USERS, MANUFACTURERS, AND INTERMEDIARIES IN IoT SECURITY

ELSA REBECA TURCIOS RODRIGUEZ

ONE THING AFTER ANOTHER

THE ROLE OF USERS, MANUFACTURERS, AND INTERMEDIARIES IN IOT SECURITY

ONE THING AFTER ANOTHER

THE ROLE OF USERS, MANUFACTURERS, AND INTERMEDIARIES IN IOT SECURITY

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus, Prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on
Tuesday 4th of July 2023 at 15:00 o'clock

by

Elsa Rebeca TURCIOS RODRÍGUEZ

Master of Science in Management of Technology, Delft University of Technology
born in Masaya, Nicaragua.

This dissertation has been approved by the promotors:

Prof. dr. M.J.G. van Eeten

Dr.ir. C. Hernández Gañán

Composition of the doctoral committee:

Rector Magnificus

Chairperson

Prof. dr. M.J.G. van Eeten

Delft University of Technology, promotor

Dr.ir. C. Hernández Gañán

Delft University of Technology, promotor

Independent members:

Dr. E. de Busser

Leiden University

Prof.dr. M. Junger

University of Twente

Prof.dr.ing. A.J. Klievink

Leiden University

Prof.dr. T. Moore

Tulsa University

Prof.dr.ir. G. Smaragdakis

Delft University of Technology

Reserve member:

Prof.mr.dr. J.A. de Bruijn

Delft University of Technology

This research has been supported by the “Netherlands Privacy Cyber Security U.S.-Netherlands Cyber Security Research Programme” with project number 628.001.033, which is partly financed by the Dutch Research Council (NWO).

Keywords: Internet of Things, Cleanup IoT malware, User experience with IoT malware, IoT malware mitigation

Printed by: Gildeprint

Cover image: Fernando Cruz Rodríguez

Copyright © 2023 by E.R. Turcios Rodríguez

ISBN: 978-94-6419-829-4

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>

To my mother

CONTENTS

Summary	1
Samenvatting	5
1 Introduction	9
1.2 Intermediaries' role in IoT security	11
1.3 Research objective	12
1.4 IoT security challenges	12
1.4.1 Evolving nature of IoT malware	13
1.4.2 Insufficient threat detection and information asymmetry.	13
1.5 Towards possible solutions	15
1.6 IoT stakeholders analysis	15
1.7 Literature review	18
1.7.1 Notifications	18
1.7.2 Users security behavior.	18
1.7.3 IoT manufacturers' identification	19
1.7.4 Botnet mitigation via intermediaries	19
1.8 Research gaps	20
1.9 Problem statement & research question and sub-questions	21
1.9.1 Study 1-3: Users studies and ISPs notifications	21
1.9.2 Study 4: Manufacturers.	22
1.9.3 Study 5: Intermediaries' prevention.	23
1.10 Dissertation outline.	23
2 User Compliance After IoT Malware Notifications	25
2.1 Introduction	26
2.2 Context	28
2.3 Related work.	30
2.4 Methodology.	32
2.5 Ethical considerations	38

2.6	Findings	39
2.6.1	Age and gender.	40
2.6.2	Device type	40
2.6.3	Comprehension.	42
2.6.4	Motivation	43
2.6.5	Compliance	43
2.6.6	Modeling compliance.	45
2.6.7	Modeling cleanup	50
2.7	Customer experience	53
2.8	Limitations and future work	53
2.9	Discussion and conclusions.	54
3	Real-World Interventions in Smart Home Security	57
3.1	Introduction	58
3.2	Background	60
3.2.1	Attacks on consumer IoT devices	60
3.2.2	Improving consumer IoT security.	60
3.3	Methodology.	61
3.3.1	Overall approach	62
3.3.2	Think-aloud protocol	63
3.3.3	Pilot	65
3.3.4	Participants.	65
3.3.5	Measuring cleanup	66
3.3.6	Ethics	66
3.3.7	Limitations.	67
3.4	Results.	67
3.4.1	Identifying suspect devices in the home.	69
3.4.2	Taking action with a suspect device.	70
3.4.3	Inferring the success of remediation	71
3.4.4	Motivation under uncertainty	72
3.4.5	The end: remediation, and reinfections	73
3.5	Discussion	74
3.5.1	Informing effective interventions	75
3.5.2	Implications for evolving IoT threats	77
3.5.3	Recommendations	78
3.6	Related work	79
3.7	Conclusion	80

4	Remediating Persistent IoT Malware	81
4.1	Introduction	82
4.2	Background	84
4.2.1	QSnatch and persistent malware	84
4.2.2	QSnatch remediation mechanisms	85
4.3	Methodology	87
4.3.1	Survival analysis	88
4.3.2	Interviews	91
4.3.3	Cleanup time after notification	94
4.3.4	Ethical considerations	94
4.4	Results	95
4.4.1	Survival analysis	95
4.4.2	Interviews	98
4.5	Related work	105
4.6	Discussion	106
4.6.1	Success and timeliness of remediation	106
4.6.2	Learning from the idealized IoT user	107
4.6.3	Self-efficacy and device compromise	107
4.6.4	Limitations and future work	108
4.6.5	Recommendations	108
4.7	Conclusion	110
5	IoT Manufacturers' role in Device Infections	113
5.1	Introduction	114
5.2	Context	116
5.3	Ethical considerations	117
5.4	Methodology	118
5.5	Findings	124
5.5.1	Manufacturers	125
5.5.2	Devices	128
5.6	Updates and security advice	129
5.7	Related work	131
5.8	Limitations and future work	134
5.9	Conclusions and discussion	135

6	Understanding Protective DNS Adoption Factors	139
6.1	Introduction	140
6.2	Methodology	142
6.2.1	Recursive DNS resolvers measurement	142
6.2.2	Prolific survey	143
6.2.3	ISP customers interviews	144
6.2.4	Enterprise interviews	145
6.2.5	Experts interviews	146
6.3	Ethics	146
6.4	Findings	147
6.4.1	Protective DNS adoption	147
6.4.2	Prolific survey	148
6.4.3	ISP customers interviews	152
6.4.4	Enterprise interviews	154
6.4.5	Experts interviews	157
6.5	Discussion	161
6.5.1	Intention vs behavior	161
6.5.2	Pros and cons of PDNS by default	162
6.5.3	Government initiatives	163
6.5.4	Recommendations	163
6.5.5	Limitations and future work	164
6.6	Conclusion	164
7	Conclusion	167
7.1	Summary of the findings	167
7.1.1	Chapter 2: User compliance after IoT malware notifications.	167
7.1.2	Chapter 3: Real-world interventions in smart home security.	168
7.1.3	Chapter 4: Remediating persistent IoT malware.	168
7.1.4	Chapter 5: IoT manufacturers' role in device infections	169
7.1.5	Chapter 6: Understanding protective DNS adoption factors	169
7.2	Discussion	170
7.2.1	Users.	170
7.2.2	Manufacturers	171
7.2.3	Intermediaries	172

7.3	Implications for governance and policy making	173
7.3.1	Hierarchical governance	174
7.3.2	Network governance	175
7.3.3	Market governance	176
7.4	Future work directions	177
Bibliography		179
Appendix A		211
A.1	Correlation between the steps performed by customers	211
A.2	Likelihood ratio test compliance models	211
A.3	Survey protocol	212
A.4	Notifications	215
Appendix B		217
B.1	Notification message and instructions.	217
B.2	Think-aloud protocol	218
Appendix C		221
C.1	Email notification content	221
C.2	Interview questions.	222
Appendix D		223
D.1	Manufacturers offering software/firmware and security advice	223
Appendix E		227
E.1	Public DNS resolvers classification	227
E.2	Focus groups and pilot	227
E.3	DNS measurement	228
E.4	Survey instrument	228
E.5	Interview protocol ISP case study.	228
E.6	Enterprise interviews	228
E.7	Experts interviews	229
E.8	Qualitative coding	229
E.9	Variables included in the final ordinal regression model	230
E.10	Top 20 countries with PDNS usage	230
E.11	Ordinal logistic regression	231
Authorship Contributions		237
List of Publications		239

Datasets	241
About the Author	243

ACRONYMS

CVEs	Common Vulnerabilities and Exposures.	13
DDoS	Distributed Denial of Service.	1, 10, 13
DNS	Domain Name System.	2, 10, 11, 12
ENISA	European Union Agency for Cybersecurity.	15, 171
ETSI	European Telecommunications Standards Institute.	15
IETF	Internet Engineering Task Force.	19
IoT	Internet of Things.	vii, 1, 2, 9, 10, 11, 12, 13, 14, 15, 16
IP	Internet Protocol.	11, 12
ISP	Internet Service Provider.	2, 22
ISPs	Internet Service Providers.	1, 2, 11, 14, 19
malware	malicious software.	1, 2, 10, 13
MoU	Memorandum of Understanding.	15
MUD	Manufacturer Usage Description.	19
NTIA	National Telecommunications and Information Administration.	15
OECD	Organisation for Economic Co-operation and Development.	9
PCs	Personal Computers.	12, 15
RFC	Request for Comments.	19
SBOM	Software Bill of Material.	15

SUMMARY

In recent years the number of Internet-connected devices (aka as Internet of Things (IoT)) has increased dramatically. IoT Manufacturers have launched into the market a variety of IoT products to make a profit, while users buy them for the convenience of the technology. Despite IoT technology's benefits to society, infected IoT devices with malicious software (malware) are a serious security concern. For instance, in 2016, we witnessed one of the largest Distributed Denial of Service (DDoS) attacks facilitated by IoT devices. This attack disrupted major well-known websites, including Twitter, Spotify, Github, and others. Infected IoT devices cause negative externalities. A negative externality is the cost that third parties, who are neither the seller nor the buyer of IoT devices, must incur to protect themselves against DDoS attacks.

In the traditional personal computers world, compromised machines can be remedied with self-service solutions like antivirus. However, there is a lack of such tools to help users remove malicious software once it has taken hold for the wide variety of IoT devices. This, in turn, creates usability issues for users in the IoT space. To remediate infected IoT devices, users may need to take different actions. These actions depend on the device type, its manufacturer, patches or software updates available, and available settings of the devices.

The IoT market also suffers from information asymmetry because users have less information than manufacturers about the security state of IoT products. More importantly, no systems are in place to alert users about security issues once they begin using the devices. Some Internet Service Providers (ISPs) (referred interchangeably as *intermediaries* in this dissertation) have undertaken the task of notifying users about infected IoT devices in their home network. These types of notifications can aid the threat detection mechanisms of infected IoT devices for users. However, not all ISPs are motivated to endeavor this task, so the ones who do, do so in a voluntary manner.

Because of the dynamic nature of this problem (which changes as attacks and IoT devices do), there is no simple solution to infected IoT devices. This is a multi-stakeholder issue in which the actors involved have different incentives. Thus, it is possible that if all parties concerned worked together, a better security outcome could be attained.

Considering that the IoT technology has certain limitations, and users will have to deal with infected IoT devices, and the aforementioned actors are involved, we set ourselves to

answer the following research question: *How can users mitigate infected IoT devices? And what role can manufacturers and intermediaries play in supporting them?* The following chapters report different studies that together address this overarching research question.

Chapters 2-4 present different user studies on how users handle infected IoT devices. In chapter 2, we study the relationship between user's compliance with some recommended steps to remediate a nonpersistent IoT malware infection and cleanup of their infected devices. Users show motivation to comply with the provided steps, and compliance with these steps increases the probability of cleaning up the infection by 32%. In chapter 3, we observe in depth the process that takes place in users' homes after receiving a notification of an infected IoT device in their home network. Users are motivated to clean up their devices even though the process takes time and effort. In chapter 4, we compare clean-up rates of persistent IoT malware versus Windows and non-persistent IoT malware. Our findings suggest that persistent IoT malware takes longer to remediate, and receiving the notification from an external party (users' ISP) played an important role in remediation.

In all user studies, we observe usability challenges that users have to face due to the generic nature of the advice they receive from their ISP. Users do what they can with the advice and tools at their disposal, hereby using a variety of familiar solutions such as disconnecting the devices to deal with the infection.

In chapter 5, we uncover manufacturers that get most often compromised in the wild. Only nine manufacturers are responsible for almost 50% of IoT infections. This highlights the importance of focusing on improving the security posture of a subset of manufacturers. Furthermore, manufacturers can have a large impact on security by removing default or easy-to-guess credentials from their development practices and set-up process of devices.

In chapter 6, we study the adoption of a protective DNS service that leverages DNS to prevent malicious network traffic in users' networks. We uncover that users are willing to pay for such a service and that most participants prefer the service if it is offered by their ISP rather than governments.

These three stakeholders, namely users, manufacturers, and ISPs, have something to contribute to mitigating infected IoT devices. The Internet as a critical infrastructure is at stake. Users are motivated to act, but they cannot if they are in the dark with no information about the security state of their IoT devices and lack actionable advice. Manufacturers, on the other hand, need to improve their security posture, and a simple action such as removing default credentials from devices' set-up process can have a large impact on IoT security. Threat detection can be facilitated through intermediaries such as ISPs by offering notifications and mechanisms to prevent infections via DNS. Additionally, governments need to play their part in incentivizing intermediaries to become more involved in this problem

due to their privileged position to mitigate infected IoT devices, as well as incentivizing manufacturer to improve their security posture. Users alone cannot mitigate infected IoT devices even if they are motivated, and the actions of manufacturers and ISPs are crucial in this endeavor.

SAMENVATTING

In de afgelopen jaren is het aantal apparaten dat met het Internet verbonden is (ook wel IoT genoemd) enorm toegenomen. Fabrikanten lanceren op de markt verschillende IoT-producten om winst te maken, terwijl gebruikers ze kopen voor de handigheid. Ondanks de voordelen die IoT technologie de samenleving kan brengen, vormen IoT-apparaten die geïnfecteerd zijn met schadelijke software een gevaar voor algemene veiligheid online. In 2016 hadden we bijvoorbeeld te maken met een van de grootste Distributed Denial of Service (DDoS) aanvallen, die mogelijk werd gemaakt door IoT-apparaten. Deze aanval ontwrichtte grote websites zoals Twitter, Spotify en Github. Aangetaste IoT-apparaten kunnen echter *negative externalities* veroorzaken. Een *negative externality* is de kosten die derden, die noch de verkoper noch de koper van iot-apparaten zijn, moeten maken om zich tegen DDoS-aanvallen te beschermen.

In de traditionele pc-wereld kunnen geïnfecteerde machines hersteld worden met persoonlijke software zoals antivirus. Er is echter een gebrek aan soortgelijke software om gebruikers te helpen om schadelijke software van hun IoT-apparaten te verwijderen. Dit creëert echter bruikbaarheidsproblemen voor gebruikers in het domein van IoT. Om geïnfecteerde IoT-apparaten te herstellen, moeten gebruikers mogelijk verschillende acties ondernemen. Deze acties zijn afhankelijk van het apparaat dat ze gebruiken, de fabrikant, beschikbare patches of software-updates en beschikbare instellingen van de apparaten.

De IoT-markt lijdt ook aan *information asymmetry* omdat gebruikers minder informatie hebben over de beveiligingsstatus van IoT-producten dan de fabrikanten. Nog belangrijker is dat er geen systemen aanwezig zijn om gebruikers te waarschuwen voor beveiligingsproblemen zodra ze de apparaten in gebruik nemen. Sommige internetproviders (ISP's) (in deze dissertatie vaak aangeduid als *intermediaries*) hebben de taak op zich genomen om gebruikers te informeren over geïnfecteerde IoT-apparaten in hun thuisnetwerk. Dit soort meldingen kan gebruikers helpen bij de detectie van bedreigingen van geïnfecteerde IoT-apparaten. Niet alle ISP's zijn echter gemotiveerd om deze taak uit te voeren, dus degenen die dat wel doen, doen dat op vrijwillige basis.

Vanwege de dynamische aard van dit probleem (dat met nieuwe aanvallen en IoT-apparaten mee verandert), is er geen eenvoudige oplossing voor geïnfecteerde IoT-apparaten. Dit is een probleem met meerdere belanghebbenden waarbij de betrokkenen verschillende

motieven hebben. Het is dus mogelijk dat als alle betrokken partijen samenwerken er een beter resultaat bereikt kan worden wat betreft beveiliging.

Gezien het feit dat IoT-technologie beperkingen heeft, dat gebruikers te maken zullen krijgen met geïnfecteerde IoT-apparaten en de bovengenoemde belanghebbenden betrokken zijn, stellen we de volgende onderzoeksvraag voor: *Hoe kunnen gebruikers geïnfecteerde IoT-apparaten herstellen? En welke rol kunnen fabrikanten en intermediaries spelen om hen hierbij te ondersteunen?* De volgende hoofdstukken beschrijven verschillende studies die samen deze overkoepelende onderzoeksvraag beantwoorden.

Hoofdstukken 2-4 presenteren verschillende gebruikersonderzoeken naar hoe gebruikers omgaan met geïnfecteerde IoT-apparaten. In hoofdstuk 2 bestuderen we de relatie tussen het naleven van de aanbevolen stappen om een niet-persistente IoT malware-infectie te herstellen en de daadwerkelijke opschoning van hun geïnfecteerde apparaten. Gebruikers tonen motivatie om te voldoen aan de verstrekte stappen en het volgen van deze stappen verhoogt bovendien de kans op geslaagde opschoning van de infectie met 32%. Hoofdstuk 3 beschrijft onze diepgaande observaties van het proces dat plaatsvindt bij gebruikers thuis na de ontvangst van een melding over de vondst van een geïnfecteerd IoT-apparaat in hun thuisnetwerk. Ondanks de lange tijd die nodig is om een apparaat van een infectie te herstellen, zijn gebruikers gemotiveerd genoeg om die taak uit te voeren. In hoofdstuk 4, vergelijken we opschoonpercentages van persistente IoT-malware met Windows-malware en niet-persistente IoT-malware. Het opschonen van persistente IoT-malware blijkt langer te duren en het ontvangen van een melding van een externe partij (de ISP van de gebruiker) speelde een belangrijke rol in de opschoning.

In alle gebruikersonderzoeken ontdekken we uitdagingen op het gebied van bruikbaarheid waarmee gebruikers worden geconfronteerd als gevolg van de generieke aard van het advies dat ze van hun ISP ontvangen. Gebruikers doen wat ze kunnen met het advies en de hulpmiddelen die ze hebben en gebruiken daarbij uiteenlopende maar vooral vertrouwde oplossingen, zoals het uitzetten van een apparaat, om met de infectie om te gaan.

In hoofdstuk 5 onthullen we welke fabrikanten het vaakst in het wild worden gecompromitteerd. Slechts negen fabrikanten zijn verantwoordelijk voor bijna 50% van de IoT-infecties. Hieruit blijkt hoe belangrijk het is zich te richten op het verbeteren van de beveiligingshouding van een subset van fabrikanten. Bovendien kunnen fabrikanten een grote impact op de beveiliging hebben door standaard of gemakkelijk te raden inloggegevens uit hun ontwikkelingspraktijken te verwijderen.

In hoofdstuk 6 bestuderen we de invoering van een *protective DNS*-dienst die DNS gebruikt om kwaadaardig netwerkverkeer in de netwerken van gebruikers te voorkomen. We ontdekken dat gebruikers bereid zijn voor een dergelijke dienst te betalen en dat de meeste

deelnemers de dienst verkiezen als deze door hun ISP wordt aangeboden in plaats van door overheden.

Deze drie belanghebbenden, namelijk gebruikers, fabrikanten en ISP's hebben iets bij te dragen aan het beperken van geïnfecteerde IoT-apparaten. Het Internet als kritieke infrastructuur staat op het spel. Gebruikers hebben genoeg motivatie om actie te ondernemen maar kunnen dat niet zonder informatie over de beveiliging van hun IoT apparaten en zonder bruikbaar advies. Fabrikanten daarentegen moeten hun veiligheidshouding verbeteren, en een eenvoudige actie zoals het verwijderen van standaard inloggegevens kan een grote impact hebben op de IoT-beveiliging. De detectie van dreigingen kan worden vergemakkelijkt via intermediaries zoals ISP's door meldingen en mechanismen aan te bieden om infecties via DNS te voorkomen. Daarnaast moeten overheden hun rol spelen door ISP's te stimuleren zich meer met dit probleem bezig te houden vanwege hun bevoorrechte positie om geïnfecteerde IoT-apparaten te mitigeren, en door fabrikanten te stimuleren hun veiligheidshouding te verbeteren. Gebruikers kunnen geïnfecteerde IoT-apparaten niet zelf opschonen en de acties van fabrikanten en ISP's zijn hierbij van cruciaal belang.

1

INTRODUCTION

The term Internet of Things (also known as IoT) was first used in 1999 by Kevin Ashton [30], a student at the Massachusetts Institute of Technology. According to the Organisation for Economic Co-operation and Development (OECD), “... (IoT) refers to the connecting of a growing number of devices and things over time to the Internet” [212]. 9.7 billion IoT devices were online by 2020 [264], and consumer IoT devices accounted for 63% of the market share [53]. It is projected that there will be three times as many IoT devices in use by 2030 as there were in 2020, with 23% of them located in Europe, 24% in the United States and Canada, and 26% in China [53].

IoT has been called the next disruptive technology of our lifetime [119]. These cyber-physical ecosystems have allowed for the digitalization and automation of society by making commonplace objects smart. Voice assistants, smart energy savings, and patient monitoring are already a reality offering comfort and convenience to users. The rapid expansion of the IoT portends exciting prospects for future businesses, consumers satisfaction, and economic expansion.

Although enthusiasm for this game-changing technology has led to a rise in the number of connected devices, it has also raised serious security issues. According to Kumar et al. [166] between 20% and 50% of the IoT devices are released by manufacturers into the wild allowing owners to access them with weak passwords, making them highly susceptible to cyberattacks. Poor security of IoT devices poses a threat not only to users’ privacy, confidentiality, and safety but also to the stability of the Internet as a crucial infrastructure.




Figure 1.1: The process of how an attacker controls an IoT botnet (composed of multiple infected IoT devices in users' networks), and uses it to launch DDoS attacks on different institutions. Also, the figure shows how the infected IoT devices scan the Internet to look for new devices to infect.

Attackers have taken advantage of weak passwords, continuous operation, and poor maintenance of IoT devices to build 'IoT botnets'[162]. An IoT botnet is a collection of infected IoT devices that can be utilized for criminal activity.

IoT botnets can deliver a variety of attacks, but most attackers use them to launch Distributed Denial of Service (DDoS) attacks against different institutions [284]. A DDoS attack is one in which devices that are compromised with malicious software (malware) are used by attackers to flood services or networks until they are no longer available [162]. For instance, in 2016, we witnessed a major DDoS attack against Dyn, a Domain Name System (DNS) provider, stemmed by infected IoT devices [23], and many popular websites such as Twitter, Spotify, Github, and others went offline [22]. Infected IoT devices impose a negative externality on third parties institutions that need to protect themselves against DDoS attacks. We refer as negative externality to the cost that these third parties, who are neither the seller nor the buyer of IoT devices, must incur as a result of infected IoT devices. [295]. On average, each DDoS attack results in \$200,000 in damages for institutions of all sizes, and many of the victims go out of business within a year after an attack [90]. DDoS attacks are not the end of the matter. In addition, IoT malware keeps evolving and extending its capabilities. Attackers can exploit the compromised IoT devices to spread the infection to other IoT devices in users' networks. Users then need to deal with remediating infected IoT devices. Figure 1.1 illustrates the process of attackers controlling IoT botnets to launch DDoS attacks and looking for new IoT devices to infect in users' networks.

Notwithstanding the negative externality that IoT technology can cause and security concerns for users, revenues from this technology were expected to surpass 1 billion Euros in countries like Germany and the United Kingdom in 2022 and 4 billion Euros in the United

States and China [53]. The success of IoT technology has a financial stake in the economy. Hence, there is a call to balance economic growth while improving IoT security.

IoT security can be seen as a wicked problem. A wicked problem is one that is hard to solve because it keeps changing and is complex [241]. What makes wicked problems unique is that they can be seen as the result of other problems. One could argue that infected IoT devices are simply caused by how quickly IoT technology is being adopted and manufacturers are rushing to enter the market making security hard to implement [162]. Others might argue that user awareness or willingness to take steps to secure their IoT devices is the problem [286]. After all, users are the ones benefiting from the technology and once they buy the product they should be responsible for its security. Another perspective is that governments are not doing enough to ensure that the IoT market works properly and that manufacturers deliver secure IoT products by design [169]. In wicked problems, different stakeholders have different interests and can render different judgments on possible solutions, but opinions vary greatly depending on their own interests in the problem [241]. IoT Manufacturers and users are at the core of this problem since they compose the supply and demand of this market, and governments have the daunting task of balancing IoT economic growth and security.

1.2. INTERMEDIARIES' ROLE IN IoT SECURITY

IoT manufacturers, users, and governments are not the only stakeholders who have a stake in IoT security. Intermediaries can also play a role. Intermediaries are stakeholders known for aggregating supply and demand and facilitating market processes, and their role has been evolving through the years [213]. Previous research identified Internet Service Providers (ISPs) (referred to interchangeably as intermediaries in this dissertation) as one of the most important actors that can fight botnets [28]. Thus, we consider their role crucial for IoT security.

To begin, the majority of compromised IoT devices (more than 80% according to Cetin et al. [55]) are found in broadband ISPs networks. Also, Cetin et al. [55] showed the impact that ISPs have reducing the number of infected IoT devices by notifying users of an infection in their home networks.

Second, ISPs are crucial to the functionality of IoT devices. ISPs often times are also Domain Name System (DNS) resolvers [62, 152] for the majority of users. Simply put DNS resolvers allow humans and internet-connected devices to communicate easily. For example, if a user wants to visit Google's domain name, he types `www.google.com` into a web browser, but his device is contacting an Internet Protocol (IP) address, which is a number (e.g. 216.239.32.10). DNS resolvers use the DNS network protocol to translate the domain names users type into IP addresses.

Similarly, in order to function, IoT devices must communicate with specific domain names. For instance, a Samsung TV will contact a Samsung domain name to update automatically (e.g. www.samsungupdates.com) or if the device is streaming certain content from Netflix, the TV has to contact www.netflix.com.

When IoT devices are infected with malware, they start contacting malicious domain names or IP addresses that are controlled by attackers instead, and these malicious requests can be detected and blocked via DNS. Figure 1.2 shows how DNS resolvers can leverage DNS to block malicious requests from any Internet-connected device (including IoT devices).

Because of their technical capabilities, expertise, and most infected devices being located in their networks, intermediaries can play a useful role in reducing the impact of IoT botnets.

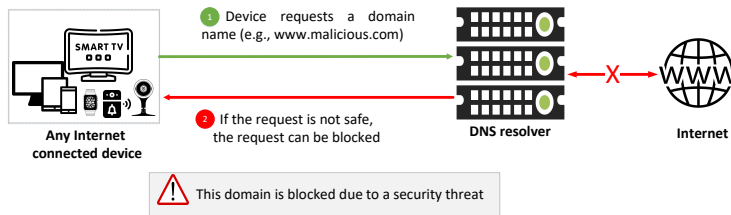


Figure 1.2: How a DNS resolver can block malicious DNS requests from Internet-connected devices.

1.3. RESEARCH OBJECTIVE

There has been a dramatic increase in the use of IoT devices in recent years. IoT manufacturers are the ones who put poor secure IoT products on the market in the first place; however, users have to deal with infected IoT devices that can be part of an IoT botnet once the devices are in use in their networks, while intermediaries are a control point that can play a role in detecting and mitigating IoT botnets. As a result, the general objective of this dissertation is to understand how users can mitigate infected IoT devices and how manufacturers and intermediaries can play a role in supporting them. Based on five peer-review studies, this dissertation aims to fulfill this goal.

1.4. IOT SECURITY CHALLENGES

Malware has been around for more than 30 years, since the first Personal Computers (PCs) with operating systems [70]. However, the proliferation of IoT devices has led to a change in malware campaigns away from traditional computers and toward IoT devices [70]. This, together with the fact that users do not have access to threat detection tools that can help

them identify infected IoT devices, are one of the most significant obstacles for IoT security.

1.4.1. EVOLVING NATURE OF IOT MALWARE

Malicious software is commonly a copy-paste of the same source code modified by attackers according to their needs [15]. Many of the IoT malware families display the same exploit capabilities, they operate by scanning open ports of IoT devices and using brute force (trying multiple easy-to-guess or common default passwords) to gain control of IoT devices [204] or use unpatched Common Vulnerabilities and Exposures (CVEs) to exploit vulnerable IoT devices [284]. Brierley et al. [47] suggests that the majority of IoT malware lack the ability to achieve persistence, as attackers often lose control of the infiltrated IoT device upon restart. A well-known non-persistent malware family is ‘Mirai’ [22]. Following certain steps such as rebooting the device and changing the default and easy-to-guess passwords remediate the infection [55].

The malware keeps evolving nevertheless, and there are already strains of persistent IoT malware. These types of malware families can maintain control of the infected device even after reboots [47]. Torii, VPNFilter [47] and QSnatch [79] are some of the most prevalent persistent IoT malware families.

Changes in IoT malware capabilities extend beyond the transition from non-persistent to persistent. In 2008, Hydra was the first DDoS-capable IoT malware to emerge; since then, several DDoS-capable IoT malware families have appeared [82]. In 2016, the largest DDoS attack in history was triggered by ‘Mirai’ malware that infected 1.2 million IoT devices [204]. However, malware targeting IoT devices can also deliver a variety of attacks [284]. This includes theft of sensitive information from infected devices, endpoint exploits (meaning any other devices in the network connected to an IoT device are compromised), crypto mining (where the IoT device’s resources are used to mine cryptocurrency), industrial spying (where industrial plants could be monitored and controlled), among others. Although users might be victims of all these types of attacks, we focus on the harm that infected IoT devices can do to third parties by means of DDoS attacks.

The growing difficulty of removing persistent IoT malware from an infected device and the evolving nature of malware capabilities, accentuates the increasing complexity and security threat that the IoT ecosystem and users have to face.

1.4.2. INSUFFICIENT THREAT DETECTION AND INFORMATION ASYMMETRY

In order to tackle infected IoT devices, it is crucial to identify the offender devices, and this is not a trivial task. The closest related work from Kumar et al. [166] successfully identifies

vulnerable IoT devices using Avast proprietary dataset, an antivirus company, and mapped them to manufacturers. However, infected IoT devices are different than vulnerable ones. At large scale, by using passive darknet data Galluscio et al. [131] identified 11 thousand exploited IoT devices. This study is a promising external measurement; however, that tells users nothing about infected IoT devices in their network. One of the challenges is to obtain the contact details of the users to communicate to them that their devices are exploited. To the best of our knowledge, only d'Estalénx and Gañán [97] have proposed a user-friendly app that consumers can simply run on their phones to detect infected IoT devices in their networks. Since this is a scholarly effort, it may take some time before it finds popular use. This is an alternative to the anti-IoT malware boxes several vendors sell to protect WIFI-connected devices. They alert about infections, and cover a limited number of devices with monthly subscriptions and they can be expensive [35].

Manufacturers may be the first to discover a security flaw in an IoT device, but they may wait to disclose this knowledge to the public until they have also produced a patch [194], leaving users in the dark about the extent to which their devices are vulnerable to a security issue. This is what is called information asymmetry. Information asymmetry is when one party has more information about a transaction than the other [11].

Nakajima et al. [194] conducted a pilot investigation with three Japanese and three American vendors and found that all of them released patches on time but did not give higher priority to addressing the most critical vulnerabilities. Guidelines exist, such as the one issued by the United Kingdom's Department for Digital Culture and Sports, which suggests that the vulnerability disclosure process should not exceed 90 days [86]. However, to the best of our knowledge, there is no extensive empirical evidence that suggests that manufacturers are following such guidelines.

A possible solution to detect infected IoT devices in users' networks is via intermediaries [55]. ISPs have the technical capabilities to detect infected customers' networks, and they can provide support and directly contact their customers. However, not all ISPs might have the incentive to perform this task. Of course, these detection mechanisms often are based on best efforts since the detection of infected IoT devices requires to have access to threat intelligence about what is malicious. This can have some cost and requires the willingness of different parties to share this information [175].

Not only does information asymmetry occur when a vulnerability or malware affects IoT devices (since users lack alerting mechanisms) or defenders have imperfect information about what they can block to prevent infected IoT devices, but also IoT device manufacturers rarely provide a thorough manual or support page and security details are typically lacking [8, 37]. This makes the 'transaction costs' [14], even for security-conscious consumers, of

buying a secure IoT device high [8, 37, 38]. Recent literature has proposed security labels [8, 101, 191], so users can make more informed choices and in turn reduce this information asymmetry. However, at the moment only Finland, Singapore, and Germany have voluntarily adopted this measure [8, 74].

When compared to the PCs industry, where consumers have long had access to antivirus software and other methods for detecting, mitigating and remediating infected PCs and informing themselves about the security of these type of devices, the Internet of Things (IoT) is still in its infancy in terms of providing similar options.

1.5. TOWARDS POSSIBLE SOLUTIONS

Best practices, standards, certifications, and legislation to address security and privacy problems related to IoT devices are being developed at a rapid pace. For instance, the European Union Agency for Cybersecurity (ENISA) has published guidelines for safe IoT software development [104]. A potential update-obligation bill is being discussed in the Netherlands, which would shift responsibility for updates from producers to retailers for IoT devices [211]. The United States' National Telecommunications and Information Administration (NTIA) drafted a legislation advocating the "Software Bill of Material (SBOM)" to improve transparency across the whole supply chain of IoT devices. The United Kingdom released a code of practice for ensuring the safety of IoT consumer devices. This code of practice aims to provide parties involved in the development, manufacturing, and retail of consumer IoT products with a set of guidelines for creating secure products from the start [87]. The European Telecommunications Standards Institute (ETSI) laid the groundwork for future IoT certification schemes and created a security baseline to include security into IoT devices from the design stage [110]. The Cyber Security Agency of Singapore and the Transport and Communications Agency of Finland have signed a Memorandum of Understanding (MoU) to recognize cybersecurity labels for IoT consumer devices in order to inform users about the cybersecurity provisions of these devices. The Federal Office for Information Security of Germany has also recently joined this effort [74].

While many of these initiatives are underway and some have already been implemented, it is very important to involve the main actors who have a stake in this problem. As noted by Rittel and Webber [241] different stakeholders may have different ideas about how to solve a wicked problem. Hence, the success of any of these best practices can be difficult to materialize if the different interests of stakeholders are not taken into account.

1.6. IOT STAKEHOLDERS ANALYSIS

Anderson and Moore [19] pointed out that security has an economic aspect, and that security

failures are caused by misaligned incentives among actors involved and information asymmetry in the market [11]. We already discussed the presence of information asymmetry in IoT security, and to understand the incentives of the actors involved in this wicked problem, first, we need to identify who are the key stakeholders. According to Perwej et al. [220] the main stakeholders in the IoT consumer market are: regulators (e.g. entities that oversee businesses, lawmakers, and certification bodies), end users, IoT device manufacturers (who produce the IoT devices and sell them), service providers (utility companies that provide hardware equipment to support or enable various smart connected home services e.g. Verisure smart alarms), platforms providers (third-party integrators like Amazon, Google), and network providers (parties who connect users' networks to the Internet, in this dissertation called intermediaries).

The objective of this research focuses on how users can mitigate infected IoT devices. However, users alone cannot remediate infected IoT devices, so the role of IoT manufacturers and intermediaries are taken into account in this dissertation. Users and IoT manufacturers are the core of the supply and demand of the IoT market and without intermediaries, it will not be possible for IoT manufacturers to reach users' homes. Furthermore, since governments are the actors who oversee end users, IoT manufacturers and intermediaries and governments pass laws to protect users' privacy and security, we consider the findings of this dissertation important to inform this stakeholder.

Each of these three stakeholders, excluding the government, has different incentives and interests in remediating infected IoT devices, but also they have a different power to contribute to solving the problem. Rooted in management theory, but also important for policy making, we use the 'Power-Interest grid' [10] to identify stakeholders' power and interest with respect to mitigating infected IoT devices. The 'Power-Interest grid' is a tool used to pay attention to and manage a specific set of stakeholders that can greatly impact achieving strategic goals and ensure its long-term viability [10]. The 'Power-Interest grid' consist of four quadrants in which actors are mapped. Stakeholders in the top two groups have the most 'interest' in solving infected IoT devices but differing degrees of power. On the top right side are stakeholders known as 'Players', they have great power to support (or sabotage) the remediation of infections. 'Subjects' (on the top left) have less influence, but they can be empowered to convert them into 'Players' to contribute to the solution of infected IoT devices.

The two lower categories are 'possible' stakeholders who haven't shown significant interest in this issue. 'Context setters' (on the lower right side) may have a lot of power, especially in terms of affecting the future context in which solutions can work, they should be encouraged to develop an interest in the problem, so that they can move up to be 'Players'.

The ‘Crowd’ (on the lower left side) has neither the interest nor the power to create an impact, but they are seen as potential stakeholders.

IoT Manufacturers are located in the ‘Players’ quadrant. IoT manufacturers have high power since they are the suppliers of these products, they are the ones who can produce patches, updates, and secure products from the design phase. Moreover, they have a high economic interest in this technology. However, investing in security, like patching devices, increase manufacturers’ costs [31]. Also, solving security issues might reduce the capacity to diversify their products or produce more devices, so they might have low interest in investing in the remediation of infected IoT devices.

Users are located in the ‘Subjects’ quadrant. They want to reap the benefits of the convenience of the technology, but they have low power either to know if a device is infected or to remediate any security issue if they do not have the support of the other actors (e.g. a security update is not available by the manufacturer to solve a security issue).

Intermediaries are in the ‘Context setter’ quadrant. They have high power due to the large number of customers they serve Internet and DNS resolutions to, and they have the contact details of their subscribers, so if they are willing, they could make users aware of infected IoT devices as was shown by Cetin et al. [55]. However, intermediaries are not part of the IoT market, so they have low interest in detecting infected IoT devices via DNS and notifying users since they might incur in costs and personnel to deal with notifications.

Figure 1.3 depicts the position of each stakeholder involved in the Power-Interest grid.

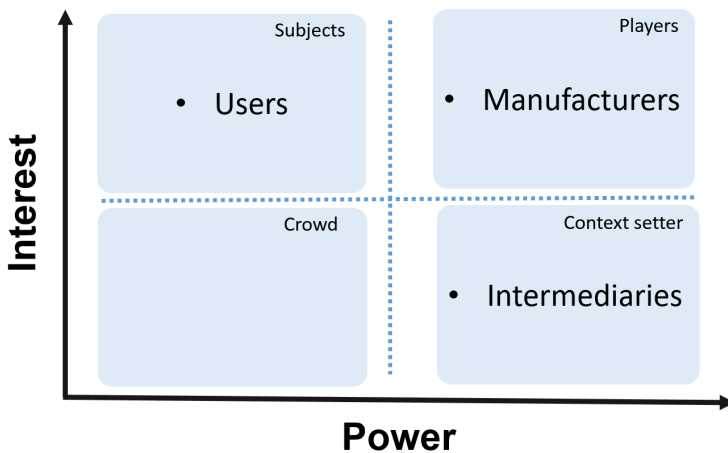


Figure 1.3: Stakeholders Power-Interest grid to mitigate infected IoT devices.

1.7. LITERATURE REVIEW

In this section, we present a review of the pertinent literature related to the current state of the art. Since this dissertation aims to understand how users can mitigate infected IoT devices and how manufacturers and intermediaries can support them, we concentrate on studies that deal with notifications' role in informing users about compromised and vulnerable devices. Next, we dig into user security behavior literature in response to security advice since users will need to act on notifications once they are notified of an infected IoT device in their network. Next, we discuss the literature on technical approaches for IoT manufacturers' identification. Finally, since most infected IoT devices can be found in intermediaries networks Cetin et al. [55], we distinguish literature investigating botnet mitigation via intermediaries.

1.7.1. NOTIFICATIONS

Notifications' role in remediating infected machines has been subject of extensive research. Li et al. [173] studied webmasters cleaning up compromised servers, and found that direct engagement with webmasters enhances cleanup by 50% and reduces infection lengths by 62%. Also, Durumeric et al. [95] observed that Heartbleed notifications led to 50% more hosts patched after being alerted of the vulnerability. Li et al. [172] also warned hundreds of network operators about security issues in their network and found a favorable impact on remediation. Stock et al. [266] notified 24,000 domains with vulnerabilities exploring different types of mechanisms of notification finding that email notifications have pitfalls, but also alternative channels of communication were not promising either. Çetin et al. [56] shows that quarantining appears to be more effective than other remediation strategies against vulnerable domains; nevertheless, this is an approach that cannot be easily implemented for this form of Internet abuse. Çetin et al. [54] have examined walled garden notifications (or quarantine notifications) and observed high remediation rates for Windows-based malware cleanup and open resolvers. Except for Cetin et al. [55] work who looks at logs of users' reactions to infected IoT devices notifications and calling users of these devices, none of the previous literature has studied users' reactions after receiving notifications regarding an infected IoT device in their network. Thus, IoT notifications after an infection is a nearly untapped topic.

1.7.2. USERS SECURITY BEHAVIOR

In a perfect scenario, where users are notified about an infected IoT device in their network, the success of the notification depends on the users' response to this notification. In light of this, it is crucial to examine user behavior in response to security advice. Research conducted

by Fagan and Khan [114] looked at what drives users to take security recommendations seriously and they found that seeing value in the advice, risk perception, convenience to execute the advice, and individual concerns motivate users to make more secure decisions. Users' reluctance to update software was investigated by Vaniea et al. [280] finding that unwillingness to update was related to three themes namely problems with figuring out if an update was "worth it," surprises with newly added functionalities, and not knowing why an update was required. Rader et al. [229] study investigated the impact of hearing about security incidents on individuals' security behavior. They found that users' security practices may be based on stories they hear that might influence their thoughts and response to security threats. To the best of our knowledge, there is no literature understanding users' behavior after an IoT malware notification. However, this is a stepping stone to involve users in the mitigation of infected IoT devices.

1.7.3. IOT MANUFACTURERS' IDENTIFICATION

As discussed in the subsection 1.4.2, to the best of our knowledge, at large scale the only studies that manage to identify IoT devices and map them to manufacturers have been published by Kumar et al. [166], and they focus on vulnerable devices. Hence, there is an open opportunity to identify which IoT manufacturers get most often compromised with malware in the wild. This can reduce the information asymmetry that users face regarding the security posture of the supplier of the IoT devices they buy.

1.7.4. BOTNET MITIGATION VIA INTERMEDIARIES

Request for Comments (RFC) 6561, recommends as best practices that ISPs notify users, so they can remediate botnets [176]. An RFC is a formal document describing specifications for a particular technology drafted by the Internet Engineering Task Force (IETF), the premier standards development organization for the Internet [148]. Furthermore, RFC8520 proposes a Manufacturer Usage Description (MUD) which aims to determine which domain names IoT devices should contact normally, so if ISPs detect abuse when devices show different behavior than the specified in the MUD, they can take action such as quarantining the device (meaning cutting down their Internet connection until the security issue is resolved) [240]. Yin et al. [300] demonstrated a method for early detection of botnets using DNS, focusing on detecting a specific type of botnet, they evaluate this solution within two ISPs. Based on the observation that IoT devices talk to particular domains, Guo and Heidemann [138] propose IOTSTEED, a system to run in routers to defend against DDoS attacks by monitoring the traffic that enters and leaves users' home networks. They suggest that ISPs should have reasons to deploy it since this could potentially protect customers from infected IoT devices

and users might be willing to pay for this service. Also, this can potentially save ISPs bandwidth that is used for illegitimate traffic instead of legitimate making their networks slow. ISPs seem to have the technical capabilities and the data to be able to leverage DNS to detect infected IoT devices, also they can play a role in notifying users [55]. Even simple ways in which these intermediaries can detect infected IoT devices is also via third parties such as Shadowserver [255] that share free threat intelligence about infected IoT devices present in different networks, so ISPs can leverage this information to inform their customers about abuse.

1.8. RESEARCH GAPS

Cetin et al. [55] demonstrates a promising approach to remediating infected IoT devices via notifications from intermediaries. While alerting users is important, the success of any notifications depends on users' behavior and actions after becoming aware of the security issue, as taking action to remove IoT malware is still in the hands of the users. Despite the fact that prior research has demonstrated the efficacy of notifications in warning users of abuse and vulnerabilities, little is known about how users react to such notifications in the IoT space. The literature suggests that users' security habits may be influenced by the stories they hear, leading to inaccurate mental models; yet, we must rely on the reality that users must learn how to react to notifications regarding security issues of IoT devices that in some cases lack even a graphical user interface.

We notice that DNS is a control point that can help keep users' networks free of infected IoT devices. When it comes to the Internet of Things, DNS methods are still in their infancy. Consequently, there is an opportunity to gain insight on how DNS might be welcomed by users as a preventative measure. Besides, although Internet service providers (ISPs) have traditionally been seen as a key control point for preventing botnets, it is unclear whether this still holds true in the Internet of Things (IoT) context and how this actor interaction with users can aid to mitigate this problem.

Furthermore, we must not forget that intermediaries are not the root cause of the problem, but IoT manufacturers are the ones who put these products on the market in the first place. We see no attempts being made to identify the IoT manufacturers that are most frequently compromised with malware and learn what measures those manufacturers are taking to address infections.

In sum, we observe three main gaps in the literature that require additional research (*i*) we do not understand how users handle infected IoT devices that have been compromised after an IoT malware notification. (*ii*) we do not have knowledge of which manufacturers get most often compromised with malware and what they are doing to solve the insecurity of

their devices. This increases the information asymmetry that users have to face regarding the security state of the IoT devices they buy. Finally, (iii) Little attention has been paid to the role of intermediaries leveraging DNS to prevent malicious activity in users' networks (including infected IoT devices), and if users are willing to accept a service leveraging DNS to protect them has not been considered in the literature.

1.9. PROBLEM STATEMENT & RESEARCH QUESTION AND SUB-QUESTIONS

Consumer IoT devices account for a significant portion of the IoT devices market share, and once infected, these devices threaten the stability of the Internet as a critical infrastructure. Users are one of the stakeholders with the least ability to solve infected IoT devices on their own. Some studies suggest [286] that users might not be willing to take steps to secure their IoT devices. However, security studies rarely prioritize investigating users' time and effort in accomplishing a security behavior; instead focusing on whether or not people actually choose the behavior that researchers suggest is desirable [248]. Using notifications as a mechanism to facilitate threat detection for users [55], there is little understanding of users' behavior after receiving such notifications in the IoT domain, which is crucial to the success of remediating IoT device infections. Also, IoT manufacturers as suppliers of this market play an important role. Understanding which IoT manufacturers are most often compromised with malicious software can reduce the information asymmetry that users have to face when buying IoT devices, but to the best of our knowledge, there is no literature covering this gap. Moreover, intermediaries can leverage DNS to prevent infections in users' networks in the first place, but little attention has been paid to their role even though the majority of compromised IoT devices are found in their networks [55]. Thus, the driving question for this dissertation will be:

How can users mitigate infected IoT devices? And what role can manufacturers and intermediaries play in supporting them?

To answer this research question, the following sub-questions need to be answered through different studies.

1.9.1. STUDY 1-3: USERS STUDIES AND ISPs NOTIFICATIONS

Building on the work on notifications, and similar to Cetin et al. [55], which to the best of our knowledge is the only study that covers end-user remediation for compromised IoT devices

after notifications, we carry out three human subject studies to observe user behavior after getting IoT notifications of a compromised IoT device by their ISP (involving intermediaries to facilitate the threat detection of malware in their IoT devices).

In the first study, after a non-persistent malware infection (specifically with Mirai malware), we contact users to study if notifications by their ISP actually lead to compliance of some recommended steps to solve the security issue, and if compliance with the recommended steps leads to remediation of the infection. Different than [55], we look at compliance as an intermediate step for the remediation of infections. Our first study answers:

RQ(s) Study1: *To what extent does IoT malware notification lead to user compliance with the indicated steps to solve the security issue? And to what extent does user compliance with those steps lead to the remediation of the infected IoT devices?*

In the second study, in order to gain an in-depth understanding of what users actually do after they are notified of a device compromised with a non-persistent malware infection (specifically with Mirai malware) and to unpack the process that occurs in users' homes that leads to the compliance with recommended steps to solve the security issue and remediation of infections (as identified in Study 1), this study addresses the following:

RQ(s) Study2: *How do end-users act on remediation advice about their infected Internet of Things device(s)?*

The final human study looks into users' experience remediating QSnatch infections, a persistent IoT malware (since IoT-persistent malware is more challenging to remediate), so this study addresses the following:

RQ(s) Study3: *Does persistent IoT malware make remediation more difficult? How do users experience their remediation effort?*

1.9.2. STUDY 4: MANUFACTURERS

Understanding which manufacturers get most often compromised can reduce information asymmetry for users about the security posture of manufacturers and have an impact on their reputation. This can serve as an incentive to improve manufacturers' security posture. Moreover, Governments, who are one of the actors whose interests are to oversee that values such as security and privacy are respected and which resources are limited, could have a clear view of which manufacturers to address to reduce the number of infected IoT devices. Thus, our fourth study provides answers to the following questions:

RQ(s) Study4: *Which manufacturers are associated with compromised IoT devices? How variable is the set of manufacturers across different countries? and What are these manufacturers doing to remediate the insecurity of their devices?*

1.9.3. STUDY 5: INTERMEDIARIES' PREVENTION

Notifications are a reactive mechanism, and they only happen once the device is already infected. Intermediaries, however, can leverage DNS to prevent malicious activity. Thus, our final study focus on users' adoption of a prevention mechanism for malicious activities (including IoT malware) that leverages DNS to prevent them. Our final study answers:

RQ(s) Study5: *What is the extent of adoption of protective DNS (PDNS) resolvers? what factors encourage or discourage the adoption of Protective DNS resolvers by users?*

1.10. DISSERTATION OUTLINE

Each chapter presents the answers to the research questions for each proposed study. Table 8 shows an overview of the different chapters in this dissertation and the peer-reviewed, empirical study that is covered.

Table 1.1: Dissertation outline

Chapter	Research question(s)	Publication
Ch.2	RQ(s) Study1	Rodríguez, E. , Verstegen, S., Noroozian, A., Inoue, D., Kasama, T., van Eeten, M., & Gañán, C. H. (2021). "User compliance and remediation success after IoT malware notifications". In <i>Journal of Cybersecurity</i> , 7(1), tyab015.
Ch.3	RQ(s) Study2	Bouwmeester, B., Rodríguez, E. , Gañán, C., van Eeten, M., & Parkin, S. (2021). "The Thing Doesn't Have a Name": Learning from Emergent Real-World Interventions in Smart Home Security. In <i>Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)</i> (pp. 493-512).
Ch.4	RQ(s) Study3	E. Rodríguez , M. Fukkink, S. Parkin, M. van Eeten & C. Gañán, "Difficult for Thee, But Not for Me: Measuring the Difficulty and User Experience of Remediating Persistent IoT Malware" <i>2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)</i> , 2022, pp. 392-409, doi: 10.1109/EuroSP53844.2022.00032
Ch.5	RQ(s) Study4	Rodríguez, E. , Noroozian, A., van Eeten, M., & Gañán, C. (2021). "Super-Spreaders: Quantifying the Role of IoT Manufacturers in Device Infections". In <i>20th Annual Workshop on the Economics of Information Security (WEIS 2021)</i> .
Ch.6	RQ(s) Study5	E. Rodríguez , R. Anghel, S. Parkin, M. van Eeten & C. Gañán, "Two Sides of the Shield: Understanding Protective DNS adoption factors" (<i>USENIX 2023</i>).

2

USER COMPLIANCE AFTER IoT MALWARE NOTIFICATIONS

ISPs are getting involved in remediating IoT infections of end users. This endeavor runs into serious usability problems. Given that it is usually unknown what kind of device is infected, they can only provide users with very generic cleanup advice, trying to cover all device types and remediation paths. Does this advice work? To what extent do users comply with the instructions? And does more compliance lead to higher cleanup rates? This study is the first to shed light on these questions. In partnership with an ISP, we designed a randomized control experiment followed up by a user survey. We randomly assigned 177 consumers affected by malware from the Mirai family to three different groups: (1) notified via a walled garden (quarantine network), (2) notified via email, and (3) no immediate notification, i.e., a control group. The notification asks the user to take five steps to remediate the infection. We conducted a phone survey with 95 of these customers based on Communication-Human Information Processing theory. We model the impact of the treatment, comprehension and motivation on the compliance rate of each customer, while controlling for differences in demographics and infected device types. We also estimate the extent to which compliance leads to successful clean-up of the infected IoT devices. While only 24% of notified users perform all five remediation steps, 92% of notified users perform at least one action. Compliance increases the probability of successful cleanup by 32%, while the presence of competing malware reduces it by 54%. We provide an empirical basis to shape ISP best practices in the fight against IoT malware.

2.1. INTRODUCTION

The number of connected Internet of Things (IoT) devices, will soon exceed the world's population [274]. On different continents, more than half of households already have at least one IoT device [166]. Although IoT is bringing convenience to people's lives, the devices also introduce serious security concerns. For a few years now, they have been compromised at scale and recruited into botnets: networks of malware-infected devices under the control of an attacker.

Many of the compromised IoT devices were put on the market without even the most basic security controls in place [145]. This puts the onus of protecting them on their users. Like with regular botnets, most compromised IoT device users are located in Internet Service Providers networks [55]. RFC6561 states that Internet Service Providers should notify users and ask them to remediate the threat [176]. Researchers [22] also argued that notifying users is an important intervention to diminish the growing number of infected devices.

A core challenge for cleanup of infected IoT is designing usable mitigation advice. Remediating infections has already been proven to be difficult for PC-based malware, where users are more likely to have workable mental models as well as effective tools, most notably anti-virus software and automatic update mechanisms. In the IoT space, the conditions for user action are much worse.

First of all, Internet Service Providers can typically not ascertain what exact device, or even what general device type, has been infected. Academic research also struggles with this problem. Antonakakis *et al.* [22] could only identify 31.5% of the 1.2M infected devices and they acknowledge that their method has an unknown error rate. Other approaches rely on intrusive traffic inspection [263] or internal network scanning [166], which are technically or legally infeasible for most Internet Service Providers. The lack of visibility into the exact device type will persist for the foreseeable future. Thus, cleanup advice has to fit, by necessity, all potential device types and remediation paths. This restricts Internet Service Providers and others to recommending a generic set of steps to the users. Each individual step may or may not be applicable and may or may not be effective in remediating the actual infection at hand.

Second, the absence of accessible user interfaces makes it difficult to perform the recommended actions or apply updates – assuming such updates are even available in the first place, which is often not the case. Combined with the lack of visibility on what device type is affected, this means that the cleanup advice cannot even tell users how to access the device to implement the required steps.

Notwithstanding these challenges, we know from recent work that providing IoT malware notifications with generic cleanup steps does in fact lead to improved remediation rates [55].

It is unknown, however, what users actually did in response to the generic and hard-to-implement instructions. No prior study has measured compliance with the recommended steps.

We present the first empirical study to measure compliance directly and improve our understanding of what users do in response to IoT malware mitigation advice. Thus, our study is able to address three key research gaps: *(i)* We do not know to what extent users comply with IoT cleanup instructions; *(ii)* We do not know if notifications cause higher compliance (compared to a control); and *(iii)* We do not know if compliance causes higher cleanup rates. The latter issue is critical in light of the grave usability problems associated with IoT cleanup advice. We cannot simply assume that trying to follow the advice actually leads to better remediation. To establish evidence-based practices in the field of IoT security, a field with growing societal impact, we need to measure two relationships. First, to what extent does user notification lead to user compliance? And second, to what extent does user compliance lead to user remediation? Prior work could not empirically estimate these relationships, because compliance has never been measured, let alone within a randomized control trial together with notification and remediation.

This paper presents a field study on self-reported user actions following an IoT malware notification. It combines a randomized control trial involving 177 customers of a broadband Internet Service Provider with a follow-up survey with 95 customers (54% response rate). We studied users' compliance with the suggested actions in the notification and how the amount of compliance affected cleanup. In sum, the contributions of this paper are:

- We present the first empirical analysis of user compliance with a notification asking them to conduct generic remediation steps for infections on any type of IoT device. We find that 92% of all notified users complied with at least one of the recommended five remediation steps. Only 24% of all notified users complied with all steps. Most users pick and choose their own path from the recommended steps. Many users also reported taking additional actions not mentioned in the notification. Even in the email-only group users comply, while they lack the incentive that quarantined users have.
- We model the impact of notifications and other predictors on user compliance and find that certain user motivations reduce compliance, while the notification comprehension did not seem to have an effect.
- We also model the impact of the amount of compliance on cleanup success. Implementing all five recommended steps increases the probability of cleanup by 32%. The notification itself has a stronger impact on cleanup than the amount of compliance. This suggests that many users chart their own course, rather than following all recom-

mended steps. We also find evidence that the presence of competing malware in the home network reduces the probability of cleanup by 54%.

- We present insights from our survey data on how consumers would like to be approached with notifications regarding IoT infections.

2

2.2. CONTEXT

Our study partners with an Internet Service Provider and its subsidiary in the Netherlands. One of the authors was embedded as an intern in the abuse department in order to conduct the study. The Internet Service Provider has been mitigating IoT infections of the Mirai family based on the abuse data it receives. We briefly discuss Mirai and then describe the notification mechanisms of the Internet Service Provider, as well as the remediation steps that the users are asked to perform.

• **Mirai Malware.** Mirai emerged in 2016 and became the malware family that demonstrated the threat posed by insecure IoT [22]. Although new families have arisen [81, 162, 286], Mirai still has a dominant presence. According to Symantec [268], Mirai was the third most common IoT threat in 2018, accounting for 16% of IoT attacks. Kaspersky mentions that Mirai families were responsible for 21% of the infected devices in that year [186]. A more recent report by IBM X-Force mentions that in the first quarter of 2019, Mirai activity doubled compared to 2018 [83]. In short, Mirai is still a relevant threat and it provides a representative case study for understanding if and how end users can perform remediation.

• **Notification Mechanisms.** Our partnering Internet Service Provider and its subsidiary brand have slightly different user populations and their own abuse handling procedures. Consumers in the subsidiary brand are notified manually on a best-effort basis, while the Internet Service Provider has an automatic procedure using abuse feeds they receive from third parties to notify consumers. Users can be notified in two ways: walled garden or email-only. (We use the term notification interchangeably with treatment. In other words, notification refers to the whole treatment that users receive.)

◦ (*Walled Garden*). This mechanism moves consumers into a quarantined network, also called a ‘walled garden’, which controls the Internet access of the users. Our partners use a so-called ‘strict’ approach, which limits all Internet access except for a set of white-listed domains [176]. Users who want to access the Internet get redirected to a landing page. The page tells them about the detected infection and instructs them to take five steps in order to solve the issue and restore Internet access. When users are quarantined, the Internet Service Provider also sends an email containing the same notification content to the user-registered contact email. Apart from notifying consumers, this process also disrupts the communication between the malware command and controls and infected IoT devices.

There are three ways by which consumers can be released from the quarantine environment. First, consumers can release themselves. To achieve this, they can fill a form explaining the steps they have taken to solve the infection and then click a button to leave the walled garden (Submitting an empty form also releases them.). The self-release option disappears if the customer suffers two subsequent quarantine events in 30 days, to avoid people releasing themselves without taking any action. The second way is to contact the Internet Service Provider's abuse department and request a release. Finally, if users did not self-release or contact the abuse department, they are automatically released after 30 days.

◦ (*Email-only.*) The second mechanism used by our partners is to warn customers only through an email. The email provides consumers with the same notification content and set of five remediation steps. The user's Internet connection remains unaffected. The main reason to use this mechanism is that the capacity of the quarantine network and the abuse department support is limited.

Examples of the notifications that consumers received are illustrated in Appendix A.4. The appendix first shows an example of the landing page users saw when they were notified via walled garden. The same content was also sent as the email notification to consumers in this treatment group. For the email-only group, the appendix shows an example notification sent to consumers in this second group. The email-only notification essentially contains the same content except that it omits statements about consumers having been placed in a quarantine environment.

• **Remediation Steps.** Both notification mechanisms, walled garden and email-only, ask the user to comply with the instruction to undertake five generic remediation steps that aim to cover as many remediation paths as possible (Figure 2.1, also see Appendix A.4). Step 1 is to identify the smart device(s) connected to the home network. The explanation mentions that likely candidates for the infected device(s) are IP cameras, Digital Video Recorders (DVRs) or similar devices, not personal computers, laptops or tablets. Step 2 is to change the password of the smart device(s). Step 3 is to restart the device(s). Since Mirai malware is not persistent, step 2 and step 3 will wipe the malware from the infected device(s) and prevent immediate reinfection via the abuse of factory-default credentials. Step 4 is to reset the modem or router to factory settings. This removes port forwarding that exposes IoT devices to the public internet, as well as a possible infection of the router itself. Finally, Step 5 is to change the password of the modem or router. These five steps are generic enough to deal with the fact that the Internet Service Provider cannot reliably identify what generic device type is infected, let alone the exact model number. Therefore, the Internet Service Provider cannot explain to the user how to exactly change the password or install an update – or even whether such an update is actually available for their device. Also, these steps are

seen as the steps that are most likely to help, but the Internet Service Provider cannot know whether they are in fact effective for each notified user.




Figure 2.1: Recommended steps in the notification mechanisms

2.3. RELATED WORK

As pointed out by [55], infected IoT devices are located in broadband networks managed by Internet Service Providers. Our research is motivated by the role that Internet Service Providers can take in notifying and warning users infected with IoT malware. Nevertheless, the warnings and notifications only work if users are able to comply and this compliance actually results in remediating the infection. We first look at the literature related to botnet mitigation by Internet Service Providers. Next, we look at work on abuse and notification and security warnings. Finally, we discuss relevant work on security behavior of users.

• Botnet Mitigation by ISPs.

Security literature highlight Internet Service Providers as a critical control point against botnets [28], and it highlights that Internet Service Providers can make a difference. Although Internet Service Providers are a critical actor to fight botnets, [26] also looked at the impact of anti-botnet initiatives on the cleanup success of botnets and concluded they have no impact. Nevertheless, they conclude that anti-botnet initiatives need to engage Internet Service Providers in taking action. Also, [222] developed a model with measures that Internet Service Providers can implement to fight botnets. They found that Internet Service Providers mainly focus on prevention and notification. In addition, RFC6561 recommends best practices that can be implemented by Internet Service Providers to notify users, so they can remediate botnets [176]. Moreover, the recent RFC8520 [167] proposes to whitelist IoT traffic through a Manufacturer Usage Description (MUD). There are discussions on how MUDs can help Internet Service Providers pinpoint abuse when devices show a different behavior than that specified in the MUD [240].

The literature has highlighted that Internet Service Providers are a relevant actor to fight botnets. Nevertheless, little attention has been paid to so far to understand whether notifying or warning users about infected IoT devices leads to user compliance with the remediation advice and whether this, in turn, leads to successful mitigation of IoT infected devices.

• Abuse and Vulnerability Notifications.

There is a large body of work on the effectiveness of abuse and vulnerability notifications by measuring the presence or absence of the security issue, without actually observing

the user's behavior. [282] studied the effect of notification content and found that verbose notifications caused more remediation of compromised websites than brief notifications. [173] studied notification content and mechanisms in terms of webmasters cleaning up compromised servers. They found that direct communication with the webmaster substantially increased the likelihood of cleanup. [59] studied the effect of the reputation of the notification sender and found that a better reputation did not improve clean up rates. [95] sent notifications for servers vulnerable for Heartbleed and found a positive impact in patching. In addition, [172] notified thousands of different network operators about security issues in their networks finding that notification has a positive impact on remediation. In contrast, [266] and [56] sent notifications to thousands of domains with vulnerabilities and found very low remediation rates. The experiment highlighted the shortcomings of email notifications and the gap between awareness of the problem and actually taking action.

Our work directly relies on two notifications mechanisms (email and walled garden) to inform infected users about the Mirai infection. Although the related work has shown that such notifications sometimes work, in certain cases low remediation rates are observed [56, 266]. A few studies looked specifically at walled garden mechanisms. [54] found high remediation rates for Windows-based malware cleanup. The only research on end-user remediation of IoT compromise, closest to our current study, found that a walled garden was also effective in cleaning up IoT malware [55]. However, these studies treat the user behavior that translates notification into remediation as a black box. No prior work has tried to observe what users actually do with the cleanup instructions, nor whether better compliance actually results in better cleanup. Our work wants to contribute to this literature by performing a real-world experiment with users that have been notified.

- **Security Warnings.** Notifications are also related to the work on security warnings. Previously, [99] studied how tolerant individuals were with delays in their activities when they are informed that they were due to security purposes, they found that users were likely to not wait when they were not properly informed that the delay was due to security purposes. Also, [100] studied web browser warnings on phishing websites manipulating the background and color of the warning to observe if users obeyed the warning. They found that text and color did not have an effect on users actually following the warning. [120] Undertook the task of designing a new SSL warning, so that they were not disregarded by users. Moreover, [12] assessed if security warnings were effective for malware and phishing websites, and they demonstrated its effectiveness in practice. Moreover, [164] looked at how users reacted to PDF download warnings and showed that these might get ignored by the user because of the exposure to false positives, incorrect mental models or not understanding that PDFs can also contain viruses. [46] used mental models to understand how advanced

and novice computer users responded to computer warnings. They reported that the groups differ in terms of how they perceive the risk they might be facing.

These studies show that user responses to notifications and warnings are highly variable and we do not yet know all the factors at play. In our study design, we incorporate factors from Communication–Human Information Processing (C-HIP) [293] theory to test if they help explaining user compliance and cleanup success. (See more details in the Methodology section 2.4).

• **User Security Behavior.** Notification mechanisms rely on consumer behavior to be effective. Fagan and Khan studied users’ motivation to follow security advice, they found that individual concern of following advice is rated higher than how this can affect others [114]. [231] looked up the immediate response of Facebook users that receive warnings about suspicious login incidents defining the common process of users to respond to the incident as consisting of incident awareness, mental model generation, and behavioral response. [280] studied how users’ negative experiences of software updates impacted their willingness to update software.

As the literature expresses user’s security practices might be based on wrong mental models, yet we need to rely on the fact that users need to learn on how to react on notifications, especially in the area of IoT devices, which present very different challenges, e.g., because of lack of a web interface on devices.

The related work has shown that notifications might work, but that their effectiveness is highly variable. The work on warnings underlines that users might ignore them. Behavioral research, moreover, has highlighted the gap between awareness and actual behavior. To the best of our knowledge no prior study, including [55], has measured compliance with IoT cleanup instructions send in a notification. As we describe in the introduction, IoT cleanup advice has huge usability problems. So, we cannot assume that following the advice actually leads to better remediation. Also, our study differs from this prior work on abuse notifications by providing the first study that opens up the black box of user behavior, most notably compliance, after receiving a notification in the area of IoT malware.

2.4. METHODOLOGY

Our data collection was carried out between May and June 2019. To answer our research questions, we combined a randomized control experiment with a survey among participants. We first randomly assigned 177 Mirai-infected customers of our partners to one of the treatments (walled garden or email-only) or to the control group. We then conducted a short phone survey based on Communication-Human Information Processing (C-HIP) theory [293]. Of the 177 participants, 95 were reachable via phone within three attempts

and accepted to respond the survey. Finally, we tracked the infections of these customers during the experiment and for two additional months, to see if the infected devices were successfully cleaned.

• **Sampling and Random Assignment.** Our partner Internet Service Provider receives a daily feed from Shadowserver containing IP addresses of Mirai-infected users in its network and that of its subsidiary brand. In collaboration with both, we used additional infection data by identifying scans that matched the Mirai fingerprint (as described by [22]) in a /15 network telescope. All identified infected users were randomly assigned to a treatment or the control group. The latter received a notification delayed by two weeks, so as to have a baseline against which to measure the impact of either notification mechanism.

Consumers detected as having infected devices during the weekends were not included in the random assignment to the treatments. This decision was made because the abuse department of the Internet Service Provider does not work during weekends. So if users needed immediate support after receiving a notification in the weekend, it would not have been possible to respond to their inquiries.

We also excluded users who had been notified about an IoT infection prior to our study. Their behavior might be different from users who were notified for the first time due to previous exposure to the remediation process. Only 9 users were excluded here.

In total, the sample consisted of 177 customers. Of these, 128 have a contract with the Internet Service Provider and 49 with the subsidiary. Our design was to randomly assign customers to three equal groups: walled garden, email-only and control. However, during the experiment we discovered that there was a malfunction with the mail server at the Internet Service Provider

Consequently, users assigned to the email-only group did not get the intended email notification at the Internet Service Provider. This meant that 43 users in the originally intended email-only group had to instead be assigned to the control group. At the subsidiary however, email notifications functioned properly. This company is smaller however, as is the number of infected users, so the email-only group consisted of 16 users. All in all, this meant that our study had a larger control group than originally intended and an email-only group that was too small to allow for strong statistical inference about its differences with the other groups. We retain the group in our analysis however for qualitative comparisons.

Table 2.1 provides an overview of the overall group assignments. It also reports the portion of each group that responded to the survey. We had high response rates in all groups.

• **Survey framework.** We used Communication–Human Information Processing (C–HIP) theory as a basis to develop our survey. To maximize the response rate, we limited the survey to require only around 10 minutes to complete. This was tested during 17 pilot interviews,

Table 2.1: Overview of group assignments and survey respondents

Group		Control	Email	Walled Garden	Total
Internet Service Provider	Participants	85	–	43	128
	Survey respondents	35 (41%)	–	28 (65%)	63 (49%)
Subsidiary	Participants	17	16	16	49
	Survey respondents	10 (59%)	11 (68%)	11 (65%)	32 (65%)
Total	Participants	102	16	59	177
	Survey respondents	45 (44%)	11 (68%)	39 (66%)	95 (54%)

which are not included in the final sample on which this study is based.

C-HIP was originally proposed as a stage model for information processing, allowing for feedback loops among stages, in which an entity tries to communicate a message to change the behavior of the receiver [293]. In our case, the Internet Service Provider and its subsidiary brand are the sources of the notification which are trying to get their customers to comply with the recommended cleanup steps¹.

We chose the C-HIP theory because it includes the source of the notification. Different sources can have different consequences on how users react. In our case, the email and walled garden seemed likely to be received quite differently. Since we had to make a trade-off between maximizing responses and the length of the survey, we study only the comprehension and motivation of the users to understand their behavior, compliance. The model includes attention, comprehension, beliefs and attitudes, and motivation. Due to the real-life settings, we could not measure the attraction that the notification caused to the users when they received it. We only notified users who were not previously notified, so this reduced the familiarity that users had regarding doing the steps and they did not have an accumulation of knowledge about the tasks. Hence, we did not measure users' attitudes and beliefs either. We cannot assume that all users comprehend the notifications, since the notifications reach users of different backgrounds with different abilities and experiences. So we have to check first whether users understand what they were asked to do. If users do not understand the notifications, they cannot correctly act upon it. In addition, motivation is key because it can activate people to comply with any directive [293]. The cost of compliance should be lower than the benefits that the users perceive by taking the recommended steps. In our theoretical framework, we also included the type of devices and demographics to control for other variables that could influence behavior that might not be related to comprehension and motivation. For instance, if the device the users' own has a web interface this could influence how easily the user can change the password of the device versus when the device

¹According to later versions of the model [69], the message needs to create an attention switch and attention maintenance in its receiver. This stage was not included in our adapted theoretical framework, since due to the real-life setting of the experiment, we were not able to measure it. Nevertheless, the notification method can trigger the users' attention

does not have a web interface. Demographics can also play a role in compliance since research has shown that characteristics such as gender and age can influence technology acceptance [267], thus how users could handle IoT devices. Hence, we want to control for these variables. This model covers two important aspects of the related work: (i) the role of the Internet Service Provider's as intermediaries and how the different types of notifications can influence compliance and cleanup in the IoT domain; and (2) drivers of user behavior, in this case comprehension and motivation, to understand the degree of compliance.

Due to the structure of sequential stages, C-HIP can be an easy tool to pinpoint where an end user drops out of the process of compliance. Each stage can be a potential bottleneck to comply. An interesting notion within the C-HIP model is that notification effectiveness can also be measured based on other stages, in this case comprehension and motivation, than the binary distinction between compliance and non-compliance. Our survey addressed the following:

Comprehension of the notification. Customers were asked if they recalled receiving the notification and if they read them. Also, they were asked if they understood the notification. The answers related to reading and understanding the notification were coded as dummy variables to measure comprehension of the notifications. *Motivation of users.* The notification must motivate customers to perform the desired behavior, in this case, to comply with the five recommended steps. Customers were asked about their motivations to comply, or not, with the notification, and their replies coded as categorical variables. After the pilot survey, these categories were refined based on the most common responses.

Compliance and additional steps. Customers were asked which steps, if any, they followed to resolve the IoT infection. The question was open-ended, so that they could tell us what they remembered doing, rather than prompting their answers by mentioning the recommended steps. We pre-coded the five steps from the notification. Steps were a binary variable that took a value of 1 if the consumer performed any of the five recommended steps or value of 0 if the consumer did not perform a step. We define compliance as the number of steps that consumers took to clean their IoT devices, out of the 5 recommended steps. As such compliance is expressed as a ratio from 0/5 steps to 5/5 steps.

Sometimes customers would mention taking additional steps, that were not mentioned in the notification. This was registered as unstructured text describing an additional step. After the survey was concluded we coded these answers in several recurring additional actions.

Figure 2.2 depicts the adapted theoretical framework used as a guide to study consumers' compliance and clean up. We also wanted to control for demographic differences and for different device types when measuring compliance and cleanup. In the survey, we also asked the type of IoT device the user thought was infected, so as to control for how device type

might impact compliance and cleanup success. Furthermore, we included demographic characteristics of the customers, recorded in the Internet Service Provider customer data, to control for differences among users. Finally, we measured cleanup success independently, to see if the self-reported compliance is predictive of remediation.

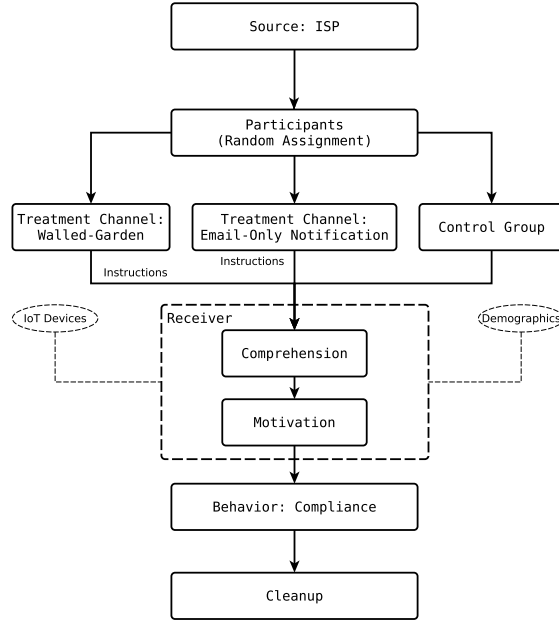


Figure 2.2: Experimental setup drawing on adaptation of C-HIP model described in [293].

Device type. Survey respondents were asked if they could identify the infected IoT device(s). Answers might be influenced by speculation or incorrect mental models. While we have no ground truth to compare these answers against, we did lookups for the customer IP addresses in Shodan [257]. Shodan is an IoT search engine, and it indexes IoT devices which means these devices are exposed to the open internet, and compared the results with the answers.

Age and Gender. We used the data of the respondent as recorded by the Internet Service Provider. When we reached someone at the listed phone number, we asked them if the subscription was in their name. In some cases, the respondent reported that they were small businesses. We coded them as such. Six users reported that someone else did the steps for them and so we coded based on their description.

Attitudes and beliefs we did not address because of time constraints in the phone call, although we did try to minimize the difference in customers' beliefs and attitudes by not including users that had been notified before.

• **Survey Process.** The survey was developed and tested in 17 pilot interviews. The pilot survey was carried out also with real consumers to check if they understood the questions,

how long the survey could take, and to refine some potential answers for the open questions. Feedback to improve the protocol was obtained and incorporated in the final design of the survey. The data of the pilot survey were not included in our results. Finally, the questions were adapted slightly, depending on if the consumer was in the control group or the treatment group. The questionnaires are included in Appendix A.3.

We conducted the survey two weeks after the notification was sent. For the control group, since they did not receive a notification, the two weeks was counted from the first day of their detection as infected. The survey call was the first notification the control group received, and for users in the control group that we could not reach by phone, we sent an email. These users, of course, were not included in the survey study, and they were not included in the measurement of the remediation rate of the control group. We set the time to contact all participants to two weeks because we want to obtain as much reliable information as possible regarding what actions a consumer took, while also giving the user time to conduct the remediation steps without being prompted to do so by the survey request. To ensure that the protocol would be consistently carried out, one person did the survey.

Survey respondents were explicitly reminded of the right to opt-out from the survey. One respondent chose to opt-out. The survey respondents did not receive any incentive to participate in the survey. Out of 177 calls we placed, 95 respondents accepted to respond the survey, one person opted-out and 81 customers could not be reached.

Because of privacy concerns, the Internet Service Provider did not allow us to record the phone survey. A script was developed to log the answers of the survey respondents. For the closed questions, the possible answers were already pre-coded. For the open questions, we added potential answers that had been given during the pilot survey and had the investigator enter manually any additional information given by the respondent.

Email logs from the abuse department were used to check if consumers contacted them for additional information. Moreover, the quarantine forms that users filled out in order to leave the walled garden were used to check if they were reporting the same device types as mentioned during the response to the survey. We used this information to validate our results.

• **Cleanup and Competing Malware.** We collected data during the experiment and for two additional months (July and August 2019) to see whether the infection was successfully removed after the experiment. We monitored the Shadowserver abuse feeds received by the Internet Service Provider [252], the Global Cyber Alliance IoT honeypot data [132], IoTPot data [216], and also a network telescope of 300K IP addresses.

We coded the infection as *cleaned* when the user's IP address was absent from the abuse reports, honeypot logs and not scanning the network telescope, either with the Mirai

fingerprint [22] or without it. We included the latter to measure cleanup conservatively. It suggests there is still an infection on the device(s) in the home network, since we would not expect a normal subscriber to aggressively scan large network blocks. We coded these cases as *no cleanup*. This analysis revealed a surprise where sometimes we found both scanning patterns for the same customer. This pattern might reveal the presence of competing malware in the home network. It has been well documented that various IoT malware families actively compete with each other for control over devices [179]. To take this factor into account, we created a dummy variable called ‘competing malware’ to capture when we saw other scanning patterns than Mirai for the same customer. To reiterate: all scanning patterns were coded as *no cleanup*.

- **One or multiple devices infected.** We were aware that the Mirai infection could be present on just one device, but also on multiple devices in the home network. Also, the ‘competing malware’ that we observed could have been present on the same device as the Mirai infection (but at a different time) or on another device. Since neither we nor the Internet Service Provider could know if one or more devices are infected, the notification was designed to handle both scenarios. It told users that one or multiple devices could be infected with Mirai. In terms of observing cleanup success, we cannot differentiate partial cleanup from no cleanup, i.e., one device was actually remediated, but another device is still infected. As long as we observed any malware scanning behavior coming from the customer IP address, we coded that case as *not clean* in order to have a conservative measurement of the remediation rate. In sum, while we lack visibility into the number of infected devices in the customer home network, we designed both the notifications as well as the measurement of cleanup to handle both scenarios.

2.5. ETHICAL CONSIDERATIONS

Our study follows the ethical principles set forth within the Menlo Report [88], namely that of respect for persons, respect for the law, justice, and beneficence. We additionally followed legal guidelines and policies set forth by our partner Internet Service Provider regarding the study and the collection of empirical data to understand consumer behavior with respect to IoT malware cleanup.

In light of the first two ethical principles (respect for persons and law), we operated within the privacy policies of our partner Internet Service Provider. One of the researchers was embedded as an intern and processed the customer data on the Internet Service Provider premise.

The survey was also conducted by the intern from within the Internet Service Provider. Consumer contact details were looked up every time prior to each phone interview and are

not part of our collected study data. All respondents were first asked for their consent to respond the survey and for the survey data to be anonymously used for the purpose of this study. The possibility to opt-out of the survey was explicitly mentioned. Only one person declined to participate in the survey. (The rest of the non-response was caused by not being able to reach the respondent.)

In terms of the latter two ethical principles (justice and beneficence), we believe that our study does not create harm and it treats individuals fairly. Our study follows a randomized control trial design (more details in §.2.4). All Internet Service Provider subscribers affected by Mirai-like malware were notified of the infection. The notification for the subscribers in the control group was delayed by 14 days. Since Mirai attacks first and foremost target third parties, not the owners of the infected devices, this delay is unlikely to expose the subscriber to substantial harm. We evaluate the downsides of this delay to be outweighed by the fact that our study aims to improve the mechanisms for users, and society at large, to combat IoT malware and prevent attacks to third parties in the future.

2.6. FINDINGS

To reiterate, our question is: to what extent do users comply with the instructions? And does more compliance lead to higher cleanup rates? We will model both relationships in light of the factors discussed in our adapted theoretical framework (§.2.4).

Table 2.2 summarizes the findings, the notifications seem to be extraordinarily effective. We calculate the odds of customers who received the notification and customers in the control group. Then we look for the odds ratio of doing one or more steps and remediation. We can observe that notified customers had 31.9 times the odds of doing more than one step than customers who were not notified. Also, we can observe that customers notified had 5.9 times the odds of successfully cleaning their infected device.

Table 2.2: Summary of findings

	No steps	One or more steps	Odds
Control	33	12	0.36
Notifications (email-only and walled garden)	4	46	11.5
	Still infected	Successfully cleaned	Odds
Control	22	23	1.04
Notifications (email-only and walled garden)	7	43	6.14

Before turning to the explanatory models, we will discuss these factors more descriptively.

2.6.1. AGE AND GENDER

To check for potential bias in the sample of participants who were reached for a survey, we compare the age and gender of survey respondents against the other participants. Table 2.3 shows the distributions. The groups are very similar across treatment conditions and demographics. Except for a bit lower proportion of female customers among the survey respondents in the control group, we see no evidence for potential bias.

Overall, the age of customers with an infected device ranges from 25 to 87 years old, with a median age of 47.5. As explained in §.2.4, when participants were reached for a survey, we asked them if the subscription was in their name. In seven cases, the survey respondents indicated it was actually owned by a small business. We coded these users separately.

Table 2.3: Study participant demographics

Group		Control		Email		Walled Garden	
		Yes	No	Yes	No	Yes	No
Survey respondent							
Age	Range	29-76	25-77	30-69	26-67	26-83	29-87
	Median	47	46	47	45	46	46
Gender	F	7 (15.5%)	14 (24.5%)	0 (0%)	0 (0%)	2 (5%)	2 (10%)
	M	31 (69%)	38 (67%)	11 (100%)	5 (100%)	32 (82%)	14 (70%)
	N/A	2 (4.5%)	2 (3.5%)	0 (0%)	0 (0%)	3 (8%)	3 (15%)
Business		5 (11%)	3 (5%)	0 (0%)	0 (0%)	2 (5%)	1 (5%)

We also compared the age of the Mirai-infected customers versus the total subscriber population of the Internet Service Provider and the subsidiary brand. We find a right-skewed distribution for the infected customers compared to the distribution of all subscribers (Figure 2.3). The mean age of Mirai-infected consumers is six years younger ($\mu=48$) than the mean age in the total subscriber population of the Internet Service Provider and the subsidiary ($\mu=54$). Welch's unequal variance t-test estimates this difference to be significant ($p = 0.0001$). In short, Mirai-infected consumers are relatively young. This fits with the speculation that younger consumers are more likely to buy IoT devices.

2.6.2. DEVICE TYPE

We asked survey respondents which of their devices they thought was infected. Table 2.4 shows the type of devices consumers reported as the offender. It is possible to notice that surveillance cameras make up a large portion of devices (36, 37.8%). This is consistent with prior studies [22, 55]. Next, 33 users mentioned a Raspberry Pi (35%). This is different from previous research. The high percentage can be understood by the fact that during our experiment, a new Mirai-based attack vector emerged targeting a known vulnerability




Figure 2.3: Age distribution - infected consumers vs all subscribers

in Domoticz software [93]. Domoticz is an open-source software that can manage home automation systems. It is often run on a Raspberry Pi. The Mirai variant exploited an ‘Unauthenticated Remote Command Execution’ vulnerability, which allowed the malware to bypass the authentication mechanism of the devices. This was detected in April 2019. Although a new version of the software was released on May 9th, 2019 [93], users reported a peak of infected IoT devices with this variant of Mirai during the study.

Nine users (9.5%) reported a Network Attached Storage device (NAS) as the culprit, which is again consistent with other studies. Next, we find a list of devices such as DVRs, Routers, Printers, Linux Embedded Systems, Smart meters, and Power consumption monitors.

Table 2.4: Infected IoT Devices

Device Type	No. Consumers
Surveillance camera	36 (37.8%)
Raspberry pi	33 (35%)
NAS	9 (9.5%)
Unknown device	8 (8.4%)
DVR	2 (2.10%)
Router	2 (2.10%)
Printer	2 (2.10%)
Linux embedded system	1 (1%)
Smart meter	1 (1%)
Power consumption monitor	1 (1%)

Surprisingly, only a small portion of the survey respondents 8 (8.4%) felt unable to identify the offending device. This could mean that most users have a pretty good understanding of their computing environment or it could mean that users are overconfident in their expertise. For example, one survey participant mentioned the ‘smart meter’ as the compromised device. The Dutch smart meters are locked-down devices that have been rolled out and maintained by the distribution grid operators. So far, there is no known attack

against these devices. Some of the answers from the survey respondents might be triggered by socially desirable behavior, as they might want to convince the investigator that they are technically savvy.

We have no ground truth against which to test the accuracy of the answers. We did conduct two crosschecks, however. First, we compared the survey answers against the submitted user forms from the walled garden. We found no inconsistencies. Second, we looked up the IP addresses of the infected IoT devices in Shodan [257]. For 36 of the 95 survey respondents (38%), we found a device listed in Shodan. Interestingly enough, 35 of these 36 (97%) survey respondents had reported the same device during the survey as was observed by Shodan. While this is hardly conclusive evidence, it does give credence to the idea that users have honestly answered our question and that they have at least a plausible speculation about the offending device. The fact that Shodan can observe it means it is exposed to the open Internet, which implies a high level of risk for poorly secured devices.

2.6.3. COMPREHENSION

In the survey, we asked whether participants received, read and understood the notification. In the walled garden group, 37 out of the 39 users (95%) remember receiving and reading the notification either via the landing page or the corresponding email. However, only 25 (67.5%) indicated they understood the notification. Interestingly, all users who acknowledged that they did not understand the message had emailed the Internet Service Provider's abuse department. In other words, even though they did not understand the notification, they all took action to find out how to solve the problem. For example, they asked for more technical information or they stated that they did not understand the cause of the infection. Of the 25 people who did claim to understand the message, 22 also emailed the Internet Service Provider. Their messages were typically stating the actions they took and then asking for confirmation whether that was enough to solve the problem.

While the email-only group was too small to make robust statements, it is worth noting that 9 of the 11 (82%) acknowledged receiving and reading the notification. Of these, 8 declared that they understood the notification. Again, the one person who did not understand emailed the Internet Service Provider's abuse department. The consumer was asking for more technical details.

In total, these results indicate that for 46 out of the 50 notified users (92%), the message was successfully delivered and read. Those recipients who did not understand the message, contacted the Internet Service Provider and asked for further details and advice. Even among people who said they did understand the message, the majority contacted the abuse department to state the actions they took.

2.6.4. MOTIVATION

We asked users an open question regarding what drove them to comply with the recommended steps. We found some recurrent topics in the answers to this question. Table 2.5 presents an overview.

In the walled garden group, 19 users (51%) said that they were driven by the fact that they did not have an internet connection. Nine users (24%) mentioned not only the lack of an internet connection, but also that safe internet is important.

In the email-only group, no one loses their Internet connection, which shifts the answers more towards more intrinsic motivations to improve security. Seven consumers in the email-only group (78%) expressed that they complied because a safe internet is important. One consumer said that a malfunctioning device was the motivation.

Table 2.5: Customer motivations to comply with notifications

Treatment	Motivation	No. Consumers
Email-only	Safe internet is important	7 (78%)
	Malfunctioning device	1 (11%)
	No answer	1 (11%)
Walled garden	Internet back	19 (51%)
	Internet back & Safe internet is important	9 (24%)
	Safe internet is important	3 (8%)
	No answer	3 (8%)
	Malfunctioning device	1 (3%)
	Need the device	1 (3%)
	Privacy concern & safe internet	1 (3%)

Similar to [114], of all notified customers only 11 (22%) expressed some social motivation to comply. Hence, it is clear that most users were thinking about how the infection affects themselves rather than others. The email-only group differs in this respect. While it is too small to draw firm conclusions, it does hint at the possibility that security practices in the IoT domain would benefit from relying on the users' social considerations regarding how infections could affect others.

2.6.5. COMPLIANCE

We asked the participants an open-ended question about compliance and then coded the answers in terms of which of the recommended steps were mentioned. We also recorded when users mentioned other steps than those recommended in the notification.

Table 2.6 displays the results for the recommended steps. Each row is one pattern of steps complied with, or not. The end of each row contains the number of users who reported this

Table 2.6: Participants self-reported compliance (1) or not (0) with each step in the notification (listed in Figure 1)

Group	Followed Steps					Freq.
	1	2	3	4	5	
Walled Garden	0	0	0	0	0	2
	1	0	0	0	0	9
	1	0	0	1	0	1
	1	0	0	1	1	4
	1	0	1	0	0	1
	1	0	1	1	0	3
	1	0	1	1	1	1
	1	1	0	0	0	2
	1	1	0	1	1	1
	1	1	1	0	0	3
	1	1	1	0	1	1
	1	1	1	1	0	2
	1	1	1	1	1	9
Email	0	0	0	0	0	2
	1	0	0	0	0	1
	1	0	0	1	0	1
	1	0	0	1	1	1
	1	1	1	0	0	1
	1	1	1	0	1	2
Control	1	1	1	1	1	3
	0	0	0	0	0	33
	1	0	0	0	0	10
	1	1	0	0	0	2

pattern. Of the 50 users who were notified and accepted to respond the survey, 12 notified users (24%) fully complied with all five steps (nine in the walled garden group and three in the email-only group). At the other extreme, 4 people in the treatment groups reported taking none of the recommended actions (two in the walled garden group and two in the email group). Taking no action whatsoever was, for obvious reasons, the dominant pattern in the control group, since they had not been notified of the problem. We will discuss this group later.

The overwhelming majority (92%) of participants in the treatment groups reported taking at least one of the recommended steps (95% of the walled garden group and 82% in the email-only group). 10% took two steps in the walled garden group and 9% in the email-only group. 26% took three steps in the walled garden group and 18% in the email-only group. 13% took four steps in the walled garden group and 18% in the email-only group. All the steps were taken in various combinations.

Even in the control group, we found that some users also reported having taken certain steps in the two weeks before, even though they had not been informed about the infection. Some users with Domoticz devices identified that their device had a security update, which

they applied. In total, ten users (22%) followed step 1. Of course, as they had not received any notification, this means that even identifying the device was a step they followed without complying with a notification. We did code it as a compliance step, to capture the degree in which users took security actions for other reasons. Two users (4.4%) followed step 1 and step 2. In the conversation with these consumers, we learned that they were prompted to take action either because of the malfunctioning of their devices or the type of device they owned.

Some combinations of steps occurred more often than others. We used Spearman's rank correlation to measure the strength and direction of the association among the steps. We observed that there is a high correlation (r_s 0.74, p 0.001) between step 2 (change the password of the device) and step 3 (restart the device). Similarly, there is a correlation (r_s 0.73, p 0.001) between step 4 and step 5: reset the modem to factory settings and change the password of the modem. This might indicate that although not all consumers did the five steps, there is some pattern to how they proceeded to mitigate the infection. Steps 2 and 3 are focused on the compromised device, while 4 and 5 are more oriented at preventing new infections. Some users focus on one, rather than the other. In Appendix A.1 the complete correlation table is presented.

We also looked at what other actions people reported, beyond the five steps. Table 2.7 summarizes the extra steps that users mentioned. As with compliance, we also include the actions taken in the control group. Interestingly, 25 (64%) of the consumers who were in the walled garden did extra steps versus 4 (44%) of the consumers in the email-only group.

Even among the users who had fully complied with the notification, some reported taking extra steps. One user, for example, described doing a software update. Among users who did not do all the steps, we found that they did report taking other actions to resolve the issue. For example, one customer reported identifying the device and also doing a software update. Other customers reported more drastic actions. After identifying the offending device, they disconnected it or stopped using it altogether. One person even mentioned that he had brought the device to the recycling center.

Of the 12 customers who took actions in the control group, some also reported extra steps. Eight customers reported doing a software update, two customers said they stopped using the device, and one customer described disconnecting the device.

2.6.6. MODELING COMPLIANCE

Almost all users in the treatment groups (92%) took some steps, though in many different combinations. Figure 2.4 shows the distribution of the count of steps taken by the users. When notified users do take action, they report on average 2.9 steps recommended by the

Table 2.7: Additional steps consumers performed

Treatment	Additional steps	No. Consumers
Email-only	Only followed notification steps	5(55.5%)
	Disconnected device	2(22.5%)
	Software update	1(11%)
	Disable port forwarding	1 (11%)
Walled garden	Only followed notification steps	12(31%)
	Disconnect device	9 (24%)
	Stop using the device	6 (16%)
	Software update	5 (13.5%)
	Disable port forwarding	3 (8%)
	Ask for help	2(5.5%)
Control group	Software update	8(18%)
	Stop use	2(4.4%)
	Disconnected device	1(2%)

notification, while the control group report on average 0.3 steps on their own initiative, without being notified.

Before turning to the models, we also did a chi-square test to validate that customers who responded were not more diligent or motivated than those who did not respond to our survey. The test result was $X^2(4, N = 177) = 0.032, p = 0.99$. The test suggests that there is no relationship between interviewed /non interviewed and clean/no clean outcomes. Also, we carried out a chi-square test to compare if the Internet Service Provider and the subsidiary had differences in the compliance steps. We only checked the walled garden group of the Internet Service Provider and subsidiary. The test result was $X^2(12, N = 39) = 5.69, p = 0.93$. The test suggests that Internet Service Provider and the subsidiary have no significant differences in terms of performed steps.

We want to understand which causal factors are associated with user compliance. We operationalized compliance as a ratio of the number of steps divided by five, the number of steps recommended in the notification (see Section 2.4). Since it is a proportion from 0 to 1, this type of data can be analyzed with a beta regression. However, beta regressions assume that the ratio is between 0 and 1, excluding the extremes. Since we also have scores of 0 and 1 in our data, we have to transform these extreme values as suggested in [73, 261]. The distribution of the dependent variable did not change.

We model the driving factors of performance using the explanatory variables from our adapted theoretical framework. Hence, we have five groups of independent variables. The first category is the treatment consumers received: walled garden notification, email-only




Figure 2.4: Distribution of the count of steps taken by the users

notification, no notification (control). Second, we include as control variables the user characteristics age, gender and status as a small business. Since there were 3 observations in the walled garden group and 2 observations in the control group with missing values for gender, we used as imputation method the most frequent value, meaning we replaced the missing values with the most common value.

Next, we control for device types. In section 2.6.2, we discussed the range of devices reported by the user. We cannot use the reported device types as explanatory factors, since many of them are used by only a few people, so the samples would be way too small to register any effect on compliance. The key difference in the population of device types consists between the Raspberry Pis and NAS and the other IoT devices like cameras and DVRs. The Raspberry Pis and NAS were specifically targeted by attackers via a known vulnerability in Domoticz (CVE-2019-10664, CVE-2019-10678). Hence, we created a categorical variable called Domoticz to distinguish between device types.

Next, we have comprehension of the notification, coded as: Understood or Did not understand. When we asked this question, there were two missing values from the walled garden group, so similar to gender, we used the most frequent value as imputation method. We ran the model with and without using the most frequent value imputation method for gender and comprehension variables, and the results did not change. And finally, we include the different motivations that were reported by users to comply. Similarly to device type, many of the motivations had a small size, so they would be way too small to register any effect on compliance. Therefore, we grouped motivations into three categories. The first category was users who wanted their internet back. The second category was composed by users whose motivations were to have the internet back and safe internet, only safe internet, and privacy concern and safe internet. Finally, other motivations include malfunctioning device, the need of the device, and no answers. Table 2.8 provides a summary of the

Table 2.8: Summary of variables

Reference category	Variables	Explanation
Control group	Walled garden	True if in walled garden group and not in email and control group
	Email-only	True if in email-only group and not in walled garden group and control group
Female	:: Age	Discrete variable
	Small business	True if it is a small business and not male and female
	Male	True if male and not small business and female
No domoticz	Domoticz	True if the device type is domoticz
Did not understand	Understood notification	True if consumer understood the notification
Internet back	Safe internet	True if motivation is not to get internet back and other motivations
	Other motivations	True if motivations are others motivation and not internet back and safe internet.

Note for Model (4) and Model (5) in table 2.9 the reference category for the walled garden group is the email-only group. For Models 1-3 N=95 and for Models 4-5 N=50. The vertical bars are to visually group the independent variables with their reference category (since Age does not have a reference category we used :: as symbol).

variables that will be included in the regression model as well as the corresponding reference categories.

Table 2.9 presents the estimated coefficient values, significance levels, and additional goodness-of-fit indicators of interest. We decided to take a stepwise approach in adding each group of variables, so we can assess their effects on compliance. Model (1) shows that the treatments – that is, the fact that users were notified – already explain 50% of the variance in compliance (R^2 0.501). Simply put: notifications do get many people to take action. This holds even for email-only. This is somewhat surprising, as earlier work [55] found that sending an email was indistinguishable from the control group, in terms of cleanup at least. In contrast, we find that emails are not ignored by users, even though they easily could do so.

From Model (4) it is also possible to observe that understanding the notification does not have a significant impact on compliance compared to consumers who did not understand since this variable only explains 6.4% of the variance in compliance (R^2 0.064). As visible in model (5), comprehension does not have a significant impact on compliance compared to consumers who did not understand the message, though the positive coefficient is in the expected direction. In terms of the different motivations, ‘other motivations’ have a significant negative impact on compliance compared to users who want their internet back. Users whose motivations are related to the need to use a device or to the malfunctioning of it, or users who did not give an answer to this question, comply less. Note that Model 4 and 5 do not include the control group, as comprehension cannot be measured for this group because they did not receive a notification. For these models, the email-only group is the reference group. Other goodness-of-fit indicators, such as log-likelihood are reported for all

Table 2.9: Estimated coefficients Beta regression on compliance ratio

<i>Dependent variable: Compliance Ratio (Transformed)</i>					
Beta Regression - link='logit'					
	(1)	(2)	(3)	(4)	(5)
Walled Garden	2.037*** (0.287)	2.036*** (0.292)	2.035*** (0.292)	0.129 (0.449)	-0.295 (0.461)
Email-only	1.928*** (0.416)	1.897*** (0.429)	1.874*** (0.433)		
Age		-0.004 (0.009)	-0.004 (0.010)	-0.007 (0.014)	-0.007 (0.013)
Small business		-0.270 (0.754)	-0.233 (0.759)	-0.769 (1.475)	-1.028 (1.443)
Male		0.008 (0.412)	-0.023 (0.421)	-0.878 (0.948)	-1.082 (0.892)
Domoticz			0.095 (0.258)	-0.044 (0.379)	0.099 (0.380)
Understood notification				0.610 (0.396)	0.340 (0.378)
Safe Internet					-0.303 (0.451)
Other motivation					-1.807*** (0.515)
Constant	-1.608*** (0.195)	-1.393** (0.624)	-1.438** (0.635)	-1.035 (1.213)	2.278* (1.201)
Observations	95	95	95	50	50
Pseudo R-squared	0.501	0.503	0.505	0.064	0.277
Log Likelihood	95.155	95.267	95.332	15.630	22.185
<i>Note:</i>				*p0.1; **p0.05; ***p0.01	

models. Higher log-likelihood values are preferred, although they alone cannot be used to determine the fit of the model. See Appendix A.2 for more details on the likelihood ratio test of the models. We will proceed to interpret our final model, model (3), as the best fit for the data.

To interpret intuitively the coefficients of model (3), the coefficients were converted to average marginal effects, Figure 2.5 presents a summary of the average marginal effects of the predictor variables on the compliance ratio, which is to say, the average expected change in compliance ratio for a change in a predictor.

We will interpret only the significant coefficients of model (3). Model (3) suggests that being in the walled garden increases the average compliance ratio by 0.39. Since the dependent variable is a proportion of the 5 steps that users took, we should multiply the coefficient 0.39. times five. Meaning that consumers in the walled garden do 1.95 steps more on average respective to the control group, which compliance ratio is on average 0,3. Similarly, receiving an email increases the average expected compliance ratio by 0.36 respective to the control group. Meaning consumers in the email group do 1.8 steps more on

average respective to the control group. Although model (5) explain less variance having other motivations rather than wanting the internet back decreases the average compliance ratio by -0.38. Meaning that consumers in the group with other motivation do 1.9 steps less on average than consumers who want their internet back.




Figure 2.5: Average marginal effect of each predictor variable

In summary, our model finds clear evidence for the impact of the notification and of user motivation. Comprehension seems to have less effect, which is somewhat puzzling, since the notifications have to rely on rather generic advice, rather than clear-cut and actionable instructions. Perhaps the generic advice is easy to understand so users do not see the subsequent questions and difficulties (‘how can I actually change the password on my IP camera of brand X?’) as a part of the message itself, but rather as a challenge separate from understanding the message. In that case, they would answer that they understood the message, even if they had trouble understanding how to comply with it.

2.6.7. MODELING CLEANUP

Now we turn to the actual goal of the notifications and compliance: cleanup of the infected devices. Figure 2.6 shows how many devices were cleaned up after two weeks of being notified or assigned to the control group, distributed over the number of steps the user reportedly took. As expected, cleanup rates are higher when the number of compliance steps increases.

An important finding is that cleanup also happens in the control group – mostly concentrated in the column with zero steps. In line with earlier work [55], we also found that 33% of the survey respondent users in the control group, who reported not taking any step, also got clean. It is unclear how this happens. We did find that around 26% of the users




Figure 2.6: Cleaned versus infected devices after 14 days

in that group also undertook action, even though they were not informed. Certain security behaviors are triggered by other mechanisms, such as update notifications. While our study added a new piece to this puzzle because users reported no action, we still cannot present a satisfying answer.

Compared to the control group, remediation rates in the two treatment groups are significantly higher. In the walled garden group, 90% got cleaned up versus 73% in the email group.

The final part of our research question is to estimate the effect of compliance on cleanup. We do this via a binomial logistic regression model. Binomial logistic regression is used when the dependent variable is binary – in this case, whether a device has been cleaned up or not.

Table 2.10 presents the estimated coefficient values, significance levels, and additional goodness-of-fit indicators. The primary focus is on the relationship between compliance and cleanup. We also look at the effect of the extra steps that consumers reported performing, at the device type, and at the issue of whether we observed scanning activity from competing malware variants for the customer. We define 3 models in which we estimate the effects of each additional variable on remediation.

An intuitive way to represent the results of binary logistic regression models is converting the coefficients into a relative risk (RR). This will capture the change of the probability of remediation after the exposure to each predictor variable. From model (1), once converting the coefficient (2.197) to RR, we can observe that an increase in the compliance ratio increases the probability of remediation by 37% as compared to the control group. In model (2) we checked the influence of device type, and it does not have a significant effect. Figure 2.7 shows the relative risk of the coefficients of our final model (3). An increase in compliance ratio increases the probability of remediation by 32%. Extra steps and the device

Table 2.10: Estimated coefficients binomial logistic regression on cleanup

Dependent variable: clean


Binomial Logistic Regression - link='logit'
(Also known as binary logistic regression)

	(1)	(2)	(3)
Compliance ratio	2.197*** (0.806)	1.524* (0.849)	1.627* (0.859)
Domoticz		0.316 (0.498)	0.648 (0.541)
Did extra step		0.869 (0.592)	0.802 (0.602)
Competing malware			-1.576*** (0.552)
Constant	0.220 (0.292)	-0.037 (0.348)	0.776 (0.474)
Observations	95	95	95
Log Likelihood	-53.782	-52.219	-47.618
Akaike Inf. Crit.	111.565	112.437	105.237
McFadden R²	0.07	0.10	0.18

Note:

*p0.1; **p0.05; ***p0.01

type have no significant effect. Competing malware presence in the home network decrease the probability of remediation by 54%.

**Figure 2.7:** Relative Risk Model 3 on remediation

2.7. CUSTOMER EXPERIENCE

The survey ended with two questions about their experience as customers of the ISP and the subsidiary brand. There was an open question asking about what consumers thought of Internet Service Provider reaching out to infected customers, and 24 (61%) of the survey respondents in the walled garden group were satisfied with this approach versus 11 (100%) in the email-only group. These results are more encouraging than in [55], in which only 17 respondents out of 76 expressed satisfaction. A possible explanation for this difference is that in our study, we asked specifically about a customer's opinion of the service, rather than analyzing the logs of people contacting the support center. The latter is likely biased towards customers being frustrated and struggling with resolving the infection. In our study, some consumers expressed frustration with losing their Internet access, but they were also glad to be contacted.

Customers were also asked for suggestions to improve the notification and remediation mechanism. Five customers in the email-only group and twenty-four in the walled garden group gave an answer. From the email-only group, two customers suggested that more information on the offender device is needed. One customer expressed that a more personalized email would help to avoid users thinking it is a phishing email. Another customer expressed the necessity of a higher availability of the abuse team, since they do not work during weekends. Finally, a customer suggested giving more publicity to the abuse team, so users would be aware of their role. From the walled garden group, twelve customers suggested that a warning prior to being in a walled garden was necessary. Along the same lines, five customers expressed that a call before putting them in a walled garden was a way to improve. Seven customers expressed that more availability of the abuse team was necessary. Other suggestions from the walled garden respondents were to explain more clearly the quarantine process and how to get out, to provide more information on the malware, to work on the authenticity of the warning, and to include information on what device type was actually infected.

2.8. LIMITATIONS AND FUTURE WORK

We discuss the main limitations of our study. First, there is the issue that there could potentially be multiple infections in the same home network. The Internet Service Provider notification did tell users that there could be more than one infection. As long as we saw signs of an infection, we coded the user as 'not clean', though the user might have cleaned up one of the infected devices. This means we cannot measure partial cleanup, only full cleanup. Second, our data is on self-reported actions. Users might have forgotten what they

did or give socially desirable answers. We cannot rule out these effects, but we did see that the devices that participants mentioned as being the culprits were, in fact, the same ones were found by Shodan at those IP addresses. We also observed that more than half of all users reported taking no action or only one step (excluding the control group, the count is one in four users). At the other extreme, only around one in four people stated doing all the steps. This pattern suggests that the tendency to provide socially desirable answers was limited.

A third limitation is the limited sample size: 177 participants in the whole study and 95 participants in the survey sample. This sample is large enough to find robust results for certain effects and causal factors. That being said, we were still left with a large portion of unobserved effects in the study on the impact of compliance on cleanup.

Finally, the experiment was carried out in one Internet Service Provider and its subsidiary brand in the Netherlands. It is unclear how well these results will generalize beyond this Internet Service Provider and country.

Future work might pursue a study with a larger sample size and in other Internet Service Providers and countries. Laboratory experiments might be an alternative, but they have their own methodological weaknesses compared to a field study with a sample of real and heterogeneous users. An important direction for future work is also to test various approaches in terms of how to actually provide usable as well as effective cleanup advice or understand why users do not take some of the suggested steps in the notification. This might need future work to collect actual ground-truth on the infected devices on customer premises, in order to have an empirical basis for remote device identification and identifying the best cleanup advice, as well as better understanding of users' mental models.

2.9. DISCUSSION AND CONCLUSIONS

Internet Service Providers are asked to implement best practices to notify consumers about IoT infections. Is cleaning IoT something that consumers can actually do? While earlier work [55] suggests that the answer is Yes, we actually knew little about the underlying mechanism. Without that understanding, we cannot design better interventions. For this reason we measured, first, whether Internet Service Provider customers complied with the cleanup advice and, second, whether this compliance improved cleanup rates.

We identified that only 24% of all survey respondents and notified participants succeeded in performing all remediation steps. The overwhelming majority of notified users, however, took at least some action upon receiving the notification. Even in the email-only group, which only received an email and had no further incentive to act, over 80% took some action. This finding suggests, differently than [56, 266], that a less intrusive notification could be

effective. However, due to the sample size more research is needed. In short, we found significant evidence that when consumers are informed about compromised IoT, they are willing to act. Users notified via email do 1.8 steps on average, while users in the walled garden do 1.95 steps on average, both compared to the control group, where users only do 0.31 steps.

When analyzing the impact on clean up, an increase in the compliance ratio increase the probability of remediation by 32%. However, if the home network was infected with competing malware, this reduced the probability of remediation by 54%. It suggests that user compliance with the recommended steps might not apply to all types of malware. Some devices remain infected or are being reinfected. IoT malware analysis has confirmed that some families fight for control over vulnerable devices. Another explanation for the effect of competing malware might be that the user owned more than one infected device. Both explanations are consistent with our finding of that competing malware are correlated with worse remediation rates.

If the impact of compliance is limited, it does not mean that the notifications as such are ineffective. Rather, it signals that the recommended remediation steps are not a sure way to get rid of the infection. Users who receive the notification might comprehend their IoT devices well enough to chart their own course out of the problem. This is supported by the fact that the impact of the notification on cleanup is higher than the impact of compliance. Cleanup rates are high in both treatment groups: 90% in the walled garden group and 73% in the email-only group. This suggests that users, once aware of the problem, are often able to resolve it, irrespective of the grave usability problems plaguing the recommended steps and IoT security in general. Putting this into the context of the C-HIP model, the notification does the attention switch that triggers users to comply. Comprehension does not play a role in changing user behavior (compliance), while the type of motivation that users expressed can negatively influence compliance compared to users who want their internet back. Its effect is not as big as the notifications. Perhaps we are seeing an effect of early IoT adopters being also more technically competent than average users. In that case, we would expect to see diminishing cleanup rates with the wider adoption of IoT.

Consistent with [114], we have observed how users' motivations are related to how the infection could affect themselves rather than how the infected IoT devices could affect others. Similar to [231], we take a step forward on understanding compliance with users fixing a real infection in their home network, giving ecological validity to these findings.

These findings clearly underline the recommended best practice for Internet Service Providers to notify infected users. Walled gardens perform the best in terms of cleanup. However, they have achieved only limited adoption among Internet Service Providers,

because of cost considerations and the fear of customer pushback. Bad luck caused our email-only group to end up too small to make robust inferences. That being said, contrary to [55], users in this group had high compliance rates and high remediation rates. Since email is a cheap and easily available option for Internet Service Providers, this could be a good second-best notification mechanism. Future work should test whether our findings for this group hold up with larger samples. In the end, though, the lion share of the burden is not borne by the Internet Service Provider. The good news from our study is that consumers are willing and able to take action, even in the absence of usable security advice and solutions.

3

REAL-WORLD INTERVENTIONS IN SMART HOME SECURITY

Many consumer Internet-of-Things (IoT) devices are, and will remain, subject to compromise, often without the owner's knowledge. Internet Service Providers (ISPs) are among the actors best-placed to coordinate the remediation of these problems. They receive infection data and can notify customers of recommended remediation actions. There is insufficient understanding of what happens in peoples' homes and businesses during attempts to remediate infected IoT devices. We coordinate with an ISP and conduct remote think-aloud observations with 17 customers who have an infected device, capturing their initial efforts to follow best-practice remediation steps. We identify real, personal consequences from wide-scale interventions which lack situated guidance for applying advice. Combining observations and thematic analysis, we synthesize the personal stories of the successes and struggles of these customers. Most participants think they were able to pinpoint the infected device; however, there were common issues such as not knowing how to comply with the recommended actions, remediations regarded as requiring excessive effort, a lack of feedback on success, and a perceived lack of support from device manufacturers. Only 4 of 17 participants were able to successfully complete all remediation steps. We provide recommendations relevant to various stakeholders, to focus where emergent interventions can be improved.

3.1. INTRODUCTION

The use of “smart” Internet-of-Things (IoT) home devices amongst consumers is growing, where this can include internet-connected home appliances, entertainment systems, and home fittings such as smart doorbells or locks. The connectivity of these devices has historically lacked sufficient security [13, 111]. Many commonly-used IoT devices have not only technical vulnerabilities, but also ineffective configuration options for password and access permissions [16, 71]. This means that a range of consumer IoT devices continue to be susceptible to malware infections, facilitating various forms of abuse, from recruiting them into botnets to personal stalking and harassment [206].

There is a direction of travel to ensure that consumers purchase secure devices, e.g., increased awareness [193], labels indicating security properties [101, 191], and improved standards of device design [42]. However, for the foreseeable future, in sufficiently secure devices continue to enter the consumer market. The brunt of the efforts to clean up infected IoT falls on both the end-users who own the devices and Internet Service Providers (ISPs), where more than 80% of the devices are located [55].

RFC6561 states that ISPs should notify users and ask them to remediate the threat [176]. Helping users protect their computer systems and remove infections has proven to be difficult for PC-based malware, even where users are more likely to have workable, effective tools available to them (for instance, automatic OS update mechanisms [290]). In the consumer IoT space, the conditions for user advice and remediation can be much more constrained when it is an ISP contacting a customer with advice; it is usually unclear what exact device, or even general device type, has been infected, forcing the advice to be highly generic. The lack of accessible user interfaces makes it difficult for users to perform the required security actions on the device they suspect is infected.

Prior work has found that notifying a user about an IoT infection can lead to cleanup [55]. Much less is known about the processes which take place in end-users’ homes after receiving a message with remediation advice. When technical experts are approached to clean a ‘smart’ personal device of suspected malware or unwanted code, they may not be able to confirm it is infected or prove removal of malware [141].

We conduct our study by partnering with an ISP which has sent notifications with remediation advice to customers infected with Mirai malware. We specifically report on the experiences of 17 ISP customers in their efforts to apply the advice. Mirai is a malware family that came to prominence in late 2016 [22], and has been referred to as the “king of IoT malware” [202]. It continues to be the leading malware family [158]. Following the notifications, we approached customers to conduct remote think-aloud observations of their attempts to follow the advice in their home, surrounded by a variety of potentially affected

devices.

We focus on the following question: *How do end-users act on remediation advice about their infected Internet of Things device(s)?* To answer this question, we documented the end-to-end story of botnet remediation which included network measurements to identify affected users, and device owner engagement. Infection data received by the ISP allows us to identify users with an infection, but also to gauge the remediation success after the intervention. We combine this with qualitative data collected during the think-aloud observations. We make the following contributions:

- We report on the real-world, in situ experiences of 17 customers acting on advice for IoT devices suspected to be infected with malware. We step out of controlled lab conditions where advice that has a known outcome is directly provided to participants. This allows us to collect data with higher ecological validity.
- We show that users are motivated, yet the advice is constrained by what can be known about the location of the infection on a home network. Many recommended actions are in practice outcomes which users must find a way to reach based on behaviours familiar to them. This adds detail to the shortfalls in the last part of advice communication for smart home users – the implications of the best-placed stakeholders (the ISP) intervening to communicate advice which is the best-available practice or which has been consolidated from manufacturers, to context-expert end-users.
- We capture the importance of advice signal design for effective behaviour change relating to smart home security hygiene. For this we relate our results to the Fogg Behaviour Model [123]. We find that where the Activation Threshold for supporting an individual to reach a target behaviour is often treated as if it were a line to cross, with home IoT it is more akin to an ‘Action Diffraction’. The user is not *able to do enough* in a direct path to the goal, due to limitations inherent in the environment, such that advocated best-practice behaviours are non-deterministic. Participants applied a range of behaviours in an approach that appeared to have a good chance of working but which were not definitely going to be successful, or be confirmed as having been successful.

The context of malware infections of consumer IoT devices is discussed in the Background (Section 3.2), including how users are typically engaged to remedy consumer IoT infections. We describe our Methodology in Section 3.3, and Results from our in situ sessions with participants in Section 3.4. The implications of our participants’ experiences are discussed in Section 3.5 and contrasted with Related Work in Section 3.6. Concluding remarks and directions for future work close the paper in Section 3.7.

3.2. BACKGROUND

Many devices enter the market that lack even basic security precautions [16]. The existence of a botnet such as Mirai starts with the manufacturing of IoT devices, which are then shipped, bought by retailers and later by consumers. Once a device has been infected, it is also unclear which of these stakeholders carries the responsibility for cleaning the device, but manufacturers generally lack incentives to prevent and remediate this problem [250].

3.2.1. ATTACKS ON CONSUMER IOT DEVICES

Different malware families use different vectors to infect vulnerable devices (such as routers, cameras and digital video recorders) [22, 55]. In the case of Mirai, there are four stages [22, 82, 156, 181, 259]. The first stage is to perform a brute-force attempt to access the device using a sequence of entries from a list of standard known username/password combinations. If this brute-force succeeds, the newly infected device sends its IP and username/password combination to the attacker. In the third stage, the report server informs the loader, which loads the malware binaries onto the device. After the binaries have been executed successfully, they are deleted, and the device is now part of the botnet.

Many IoT devices do not support standard user interfaces, which makes it difficult for customers to change the standard passwords (assuming a device has such a feature to begin with, which may not be the case [117]). Even where a device has an adequate interface, many users prefer having a working device as soon as possible over going through security-related installation steps (such as replacing the standard password) thoroughly [162] (where the inter-connected nature of smart homes means this may include securing the entire home network). End-users who do care about security may lack knowledge to perform the right actions, due to the heterogeneity of IoT devices [21, 306].

3.2.2. IMPROVING CONSUMER IOT SECURITY

Information about the security qualities of IoT devices can potentially be difficult to find. One avenue of research focuses on supporting consumers to make informed choices about the smart home devices they buy in the first instance (e.g., security labels [101, 191] and consumer guides [193]) Another area of focus has been to ensure that device design matches user needs; this has been noted regarding specific requirements for access control [142] and privacy in a shared environment [305], for instance.

Most vendors of IoT devices do not deliver a comprehensive manual or support page with their product. Where information is provided, details relating to security are often absent or not adequate [37, 128]. This means that even for those consumers who do care about security [38, 205, 246], the ‘transaction costs’ of ensuring purchase of the most secure

device are simply too high [14, 37].

As the Internet increasingly connects end-users and their devices globally, it becomes complex for governments across the world to organise clear responsibilities and liabilities for security. As the IoT is still relatively new and evolving, it could take some time before governments are able to clean the market of insufficiently secure devices and exert pressure on responsible parties. Simple improvements such as labelling the level of security of devices could improve the purchasing environment [155], but even for such small improvements, incentives are lacking. As present, the most viable mitigation techniques mostly come from Internet Service Providers (ISPs) intervening when customers' devices are compromised, or information campaigns to realise prevention through consumer awareness. However, levels of remediation are far from perfect. The content of a notification should be understandable and clear for target users, but there is a balance to be struck. Research has found that detailed steps can strengthen the effect of the notification [96, 174, 283]. On the other hand, messages should be plain and simple [126].

Even where users are aware of a security problem and activated to act, there can be uncertainty about which device is infected, or how to take the required action [233]. Users may instead rely on familiar techniques to solve problems on 'unfamiliar' devices, which often is not the correct approach for new types of devices and infections [304].

For structuring interventions, identifying critical points in life cycle of devices is useful [163]. *Opportune moments* for intervention then emerge [123], which are important for focusing resources toward enacting a behaviour at a specific point where it is more viable. Where purchase of new devices is one such point [101, 217], the notification to a customer of a suspected malware infection is another opportune moment. However, There are challenges inherent to deploying behaviour interventions where the 'influencer' does not manage the environment. In managed environments (including the artificial/controlled environment of a lab study), the influencer can know who the target is and how to reach them. Here, we study an environment where that knowledge is not immediately available. We then leverage technical tools to approximate where the intervention is needed, by triangulating across datasets to identify devices which are vulnerable. Simply put, we have to find a way to go to the participant, whereas normally in a study the participant comes to us.

3.3. METHODOLOGY

In this section we describe our approach to answering the main research question. This involved partnering with an Internet Service Provider (ISP) and studying customer responses to remediation instructions.

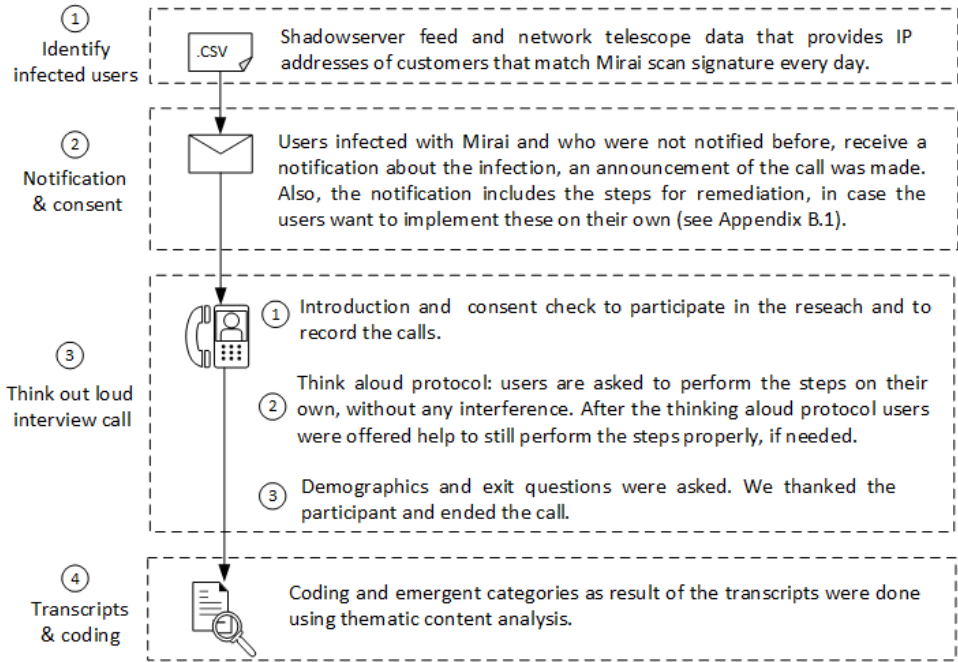


Figure 3.1: Approach and data collection.

3.3.1. OVERALL APPROACH

Our study starts with identifying ISP customers who suffer from an active Mirai malware infection. For this, we used two data sources. One was the Shadowserver drone report [253]. The ISP receives from Shadowserver a daily list of IP addresses of customers that match the Mirai fingerprint. Mirai scans have a particular signature, where an artefact of the malware's stateless scanning approach is that each probe includes a TCP sequence number equal to the destination IP address that the malware is targeting to attack [22]. This is conventionally used to detect the malware.

A network telescope was then employed. This is a set of unused IP addresses [190], where the traffic targeting this IP set is usually unsolicited. The network telescope of 300K IP addresses logs the IP addresses of hosts that were scanning with the Mirai fingerprint, as described in [22].

This is Phase 1 in the overall approach (as in Figure 3.1). The ISP is in a unique position to know which customer is associated with an IP address, so that we could identify which customers were suffering from a Mirai infection.

If the identified owner had not yet been notified, the ISP would notify the user about the infection via email (Phase 2, Figure 3.1). Included in this email would be an explanation

of the research, and an invitation to participate in a call to understand better the process that users follow to execute the steps, as part of the standard service. It is also mentioned that users are free to execute the steps themselves (see section B.1 for more details on the notification) without opting in to the study. During the call, each customer was asked explicit consent to participate in the research and record the call (see section B.2). Minimal data of customers who did not consent to be part of this research was received in advance to be able to contact the customer, but it was not included in the results of this research.

To further ensure that the email notification could be understood by those end-users who received it, several communication experts from the communication department of the ISP transcribed the text to B1 level of the Common European Framework of Reference of Languages (CEFR) [113]. This is an international standard to describe language proficiency, in which B1 indicates basic level. The email notification was written in both English and Dutch (as the main language where the study was carried out).

A day after the email notification, users would be called (Phase 3, Figure 3.1). Three users did not answer during three attempts to call them and were left out of the study. Our protocol has a check at the beginning to ensure we talk to the device owner. We then asked users whether they wanted to opt into participating in the study, asked for explicit consent to record the interview, and explained that the participant could end the call at any moment (Appendix B.2, part 1).

After concluding a call, a transcript was created. We used thematic analysis (Phase 4, Figure 3.1) to code transcribed copies of the interactions (from audio recordings). For performing the thematic analysis, the step-wise approach listed by [18] is used. Two of the researchers coded the transcripts to identify themes. Dedicated code review discussions took place between coders (to address emergent themes and conflicts), which happened in stages before arriving at the final set of themes. A balance in themes was found through iterative merging and splitting existing themes until convergence was reached into the most important themes (where the subsection in our Results represent theme families, Section 3.4). Saturation of themes was reached after 17 calls.

3.3.2. THINK-ALOUD PROTOCOL

Originally we had planned to visit customers' homes/premises, to interact with them in an a natural and comfortable environment, and be physically present when users execute the recommended remediation advice. There was a need to instead develop a novel phone-based protocol for interacting with the customers of the partner ISP, foremost due to social distancing measures (Section 3.3.6). A positive aspect of this was that all participants were at the appropriate location when they were contacted.

To prepare, experience was gained in managing cases where remediation was not possible. One of the researchers accompanied a senior mechanic from the ISP for a day, and gained insights from the ISP customer support staff regarding how to build trust with customers. In cases where the engineer is not successful in helping users, the most important step was seen as informing the consumer of the situation and to let them know about the possible ways forward. In such cases, also a supervisor should be informed about the issue. It can reach a point where informing the customer of an issue is the best one can do. This reflects the reality that the ISP is not technically responsible for the device, even though it has the opportunity to intervene.

The think-aloud protocol (Phase 3, Figure 3.1) consisted of three stages:

- **Stage 1:** Consent and notification: First, we obtained consent to conduct the study, asking then for approval to record the interview. Next, we checked whether participants received the notification and, if not, we sent it again and provided the participant time to read it.
- **Stage 2:** Acting on the advice: We allowed the participants opportunity to perform the actions and verbalize their thoughts, without direct input from the researcher. This think-aloud activity was transcribed and analysed.
- **Stage 3:** Demographics and support: We collected demographics and, if the researcher saw an action during Stage 2 as incomplete or incorrect, suggestions were offered for performing actions correctly, to the extent that this was possible (see 3.3.7). Last, we thanked the customer for their participation as well as provide e-mail details for future contact with the researcher in case they had any questions.

See Appendix B.2 for complete details on the think-aloud protocol. The technical advice provided to customers (in the email and in the second step of the protocol) are steps used by the partner ISP, so it is what the ISP considered best advice. For comparison/reference, these steps are comparable to what is advised in online sources, as found on the Krebs on Security blog¹ and Symantec/Norton website².

During a call with a participant, they would try to implement the 5 recommended actions from the email: (1) determine which devices are connected to the internet that could potentially be infected with Mirai; (2) change the password of these devices; (3) restart the devices by turning them off and on; (4) reset the modem/router to the factory settings, and; (5) change the password of the modem/router (Appendix B.1 contains the message in full).

¹<https://krebsonsecurity.com/2018/01/some-basic-rules-for-securing-your-iot-stuff/>

²<https://us.norton.com/internetsecurity-iot-smart-home-security-core.html>

Table 3.1: Summary of participants demographics, devices, actions, and outcomes. No. of users refers to the number of people in the household of the participant. Some connections were part of a small business rather than a home. Steps 1-5 refer respectively to actions relating to Device Identification, Device Password, Device Reset, Router Reset, and Router Password. Boxes highlighted in gray refer to an outcome classed as a failure to complete the associated Step, otherwise the action was a variation on a successful outcome. The letter-specific codes for each step are detailed in Figure 3.2

Index	Age	Gender	No. Users	Suspected device	Step 1	Step 2	Step 3	Step 4	Step 5	Remediated?	Reinfection?
1	53	M	6	Router	1B	n.a.	n.a.	4A	5A	Yes	No
2	55	F	1	IP camera	1A	2D	3A	4C	5C	Yes	No
3	43	M	2	IP camera	1A	2D	3A	4C	5A	Yes	No
4	49	M	3	IP camera	1A	2D	3A	4C	5A	No	Yes
5	65	M	2	IP camera	1A	2C	3A	4D	5D	Yes	No
6	21	M	Business	IP camera	1A	2B	3C	4C	5A	Yes	No
7	45	M	4	Router	1B	n.a.	n.a.	4C	5C	Yes	No
8	65	M	2	NAS	1A	2C	3A	4C	5A	No	Yes
9	61	M	2	Smart printer	1A	2C	3A	4C	5A	Yes	Yes
10	34	M	Business	IP camera	1A	2A	3B	4A	5A	Yes	No
11	55	M	Business	NAS	1A	2A	3A	4A	5A	Yes	No
12	80	M	2	Doorbell	1A	2A	3A	4C	5A	Yes	Yes
13	49	M	1	IP camera	1A	2D	3A	4A	5A	Yes	No
14	43	M	2	-	1C	2E	3D	4A	5A	Yes	Yes
15	53	M	5	Router	1B	n.a.	n.a.	4B	5B	Yes	No
16	41	M	3	IP camera	1A	2B	3C	4C	5A	Yes	No
17	42	M	4	Smart TV	1A	2C	3A	4A	5A	No	No

3.3.3. PILOT

The study protocol was tested with 7 customers. These pilot sessions were especially important for refining the protocol, as the main study would also involve interacting with real customers of the ISP and an intervention that has not been studied directly in a real-world setting. We could also evaluate the think-aloud protocol, accounting for not being present in the room with the users.

Similar to the insights from the ISP customer support staff, trust was found to be important: 5 of 7 customers were cautious about the call, 4 wanted a more detailed explanation of the research, and one called back to the service desk to confirm the authenticity of the research and email.

The pilot resulted in a check being added at the beginning of the protocol to talk to the person who takes care of security issues (as pilots included cases where the person who set up the devices did not live in the household); issues of delegation to informal technical support are discussed in [223]. The most significant change in the protocol was the inclusion of more upfront information about the purpose of both the call and research, to bolster trust.

3.3.4. PARTICIPANTS

All customers with a diagnosed Mirai infection in the period between May and July 2020 were notified by email about the infection and the study. If they did not opt out of the ISP's support process, they were called the next day. During the experiment period, 37 unique IP

addresses corresponded to 37 customers with Mirai infections. 12 were observed during the weekend, where the helpdesk at the ISP does not notify these users as they cannot provide support over the weekend. Of the 25 remaining IP addresses, 3 could not be notified due to technical issues within the ISP, 2 did not respond to attempts to contact them after being notified, and 3 were not willing to take part in the experiment (did not opt-in to the study). There were think-aloud observations with 17 customers. The age of the participants was between 21 and 80 years old with a median age of 49. We interviewed 16 males and 1 female, and from the 17 participants, 3 used their internet connection to run their own businesses. Table 3.1 shows the participants' demographics. As was also the case during the study pilot, sessions each took approximately 30 minutes in total (15 minutes of which was the think-aloud protocol).

No incentive was provided to users to participate, beyond the possibility of providing the technical support detailed in the participant-facing study materials (see section B.2).

3.3.5. MEASURING CLEANUP

From the two data sources described in subsection 3.3.1, we received daily lists of IP addresses where infected Mirai hosts were located. This led to the initial identification of the customers and the recruitment of participants. We kept monitoring this data for an additional two weeks after the call.

Mirai reinfection can occur within a few minutes, or for some devices within 48 hours [55]. We chose a conservative 4-day window to determine remediation. Since Mirai attacks involve aggressively scanning the IP space for devices, we presumed a two-week window to measure reinfections as related to the state of participants' home network. We illustrate this way of measuring outcomes in Table 3.1. We should note that this observation method is not perfect. While false positives are highly unlikely, because of the specific Mirai fingerprint, false negatives might occur (an infected host might not show up in the data, even though it is still infected).

3.3.6. ETHICS

The study protocol was approved by our institution's human research ethics committee (TPM project 1083). The study design followed the principles for ICT human research as detailed in the Menlo Report [89] (as indicated also in the design of the think-aloud protocol). To make sure the end-users feel that they are in a safe environment, the think-aloud protocol is built around ensuring that the participant feels they are in a safe space and have not done anything wrong, and can state their feelings and actions without any judgment.

The first part of the call is about informed consent. This consent involves both taking

part in this research anonymously, as well as the call taking place and the recording of it. Users were reminded that they could stop the study at any time. If they did not wish to participate, they were informed that they would be processed as usual by the partner ISP.

3.3.7. LIMITATIONS

In adherence with national social distancing measures related to the Covid-19 pandemic, in-person data collection was avoided. In-person home visits may have allowed for opportunistic observation of relevant details outside of our protocol, or differences between stated and actual behaviour. We compensated for this with a think-aloud protocol. We cannot rule out, however, that users may not have accurately described what they did via the call. Even though the researcher is trying to stay at the side-line, their presence influences the participants [154, 171, 278], who will typically pay more attention and effort to the tasks within the study. This does not detract from the context of the interaction, which would naturally require the individual to focus on the instructions regardless.

The research may have engaged with device owners who were unable to knowingly secure their devices. In such cases, at the end of the protocol they were helped to execute the steps they missed properly (after the think-aloud protocol). Also, an e-mail for future questions or contact was provided. The researcher helped the participants with any unsuccessful steps in accordance with the study protocol. Although infections could have plausibly been remediated, participants were carrying out actions themselves within the online ‘interview call’ format, and outcomes were based on customers’ reported actions. For instance, users may have changed passwords though we may not have been able to corroborate the outcome, or whether the advice absolutely caused the outcome.

Our work is based on users’ data from a single ISP. Hence, more research will be necessary to validate these results across multiple ISPs and different countries. Similarly, we focus our design and analysis on a single malware family, Mirai. The recommended steps might differ from those for other malware families. We see trends of advice only becoming more complicated, see Section 3.5.2).

A final point is that our measurements of remediation and reinfection is not perfect. The infection data suffers from a small rate of false negatives. We compensate for this by working with longer time windows. Only when participant’s IP addresses are not seen in the infection data for four consecutive days, do we conclude they successfully remediated.

3.4. RESULTS

Participant sessions were transcribed and analyzed to understand the ‘journey’ of remediation, following the steps of advice. We present our findings by following this journey. No

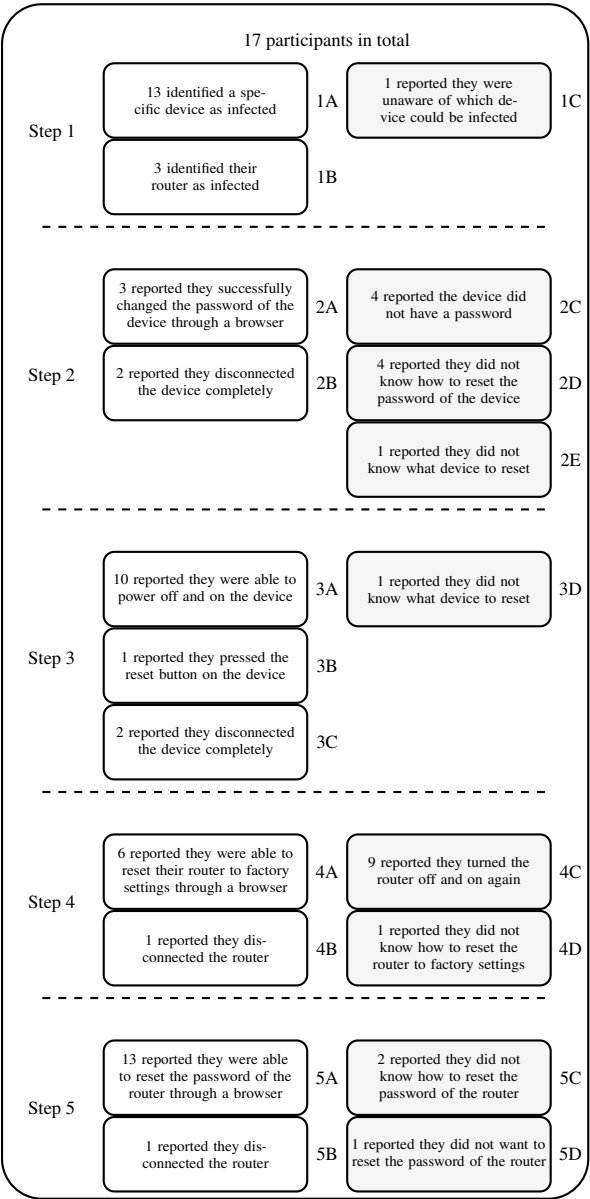


Figure 3.2: Overview of outcomes of actions by participants, while attempting to execute the remediation advice. Steps correspond to those found in Appendix B.1

participants reported having attempted to apply the steps before the session. We describe how participants attempted to: first, identify the infected device (Step 1, subsection 3.4.1);

implement the recommended actions on that device and on their router (Step 2-5, subsection 3.4.2); infer the success of their actions (subsection 3.4.3), including their motivation to work through what transpired to be an arduous process for almost all participants (subsection 3.4.4). Finally, we connect the customer experiences with our measurement data on whether the infection was remediated (subsection 3.4.5).

Figure 3.2 provides an overview of reported participant actions. Each labelled box represents a particular action. To illustrate: 13 users took action 1A and identified a specific device as infected. White boxes indicate a successful action in terms of enacting advice, grey indicates no success.

3.4.1. IDENTIFYING SUSPECT DEVICES IN THE HOME

The first remediation action is to identify which devices are connected to the internet and could be infected with Mirai. The notification email informed participants that Mirai would not be present on a regular PC, laptop, tablet or phone. The subsequent actions (changing the password and turning the device off and on) are meant to be applied to all the devices that could potentially contain Mirai. A cautious approach is then to remediate and secure all potential victim devices.

Thinking aloud, four participants immediately focused on the device that they thought was the most likely culprit. All other participants started enumerating their devices, e.g., P12: *“I have 22 devices connected to the internet. Cameras, a garden sprinkler, a doorbell, the list goes on.”*

Whether multiple devices were enumerated or not, all participants focused on identifying one suspect – no participant ended up identifying multiple suspect devices. We observed participants using three heuristics to reason about the likely culprit. The first heuristic, used by the majority of participants, was a process of elimination, as with P04: *“I have a laptop, two mobile phones, no three mobile phones. I have a camera, a security cam, and the solar energy is also connected to the internet. I run anti-virus on everything. I just bought that for five devices, also for my wife’s iPad. According to that email, it would have to be the security camera.”*

This first heuristic might not lead to a confident identification, as seen with P01: *“OK, in the email you write that it can’t be phones, laptops, or really anything with Android on it. That leaves us with printers and cameras and the like. But I don’t have those. Yeah, I have a printer, one of those all-in-one types, but that isn’t even switched on at the moment [...] So that doesn’t make sense.”*

The second heuristic, used by eight participants, was honing in on a device that the person recently experienced problems with. This occurred for instance with P02: *“I think*

it is the camera. [...] It says there is a system error and it needs a restart. But only the company can do this remotely.”, and P06: *“There are 4 phones connected to the wifi and a computer. And the security camera, but that doesn’t work properly anymore. It actually seems likely that this camera is misbehaving.”*

A third heuristic was only employed by one person: conducting an Internet search. P15: *“I have one all-in-one printer, that is never turned on, a beamer connected to the internet, an Xbox, Nintendo Switch, a smart TV, 2 laptops with Windows 10, a laptop with Windows 8 and a [routerModel] [...] Now, I saw in the email that it can’t really be one of these devices, so I searched on Google for all my devices [...] then I found that [routerModel] has been having problems in [another country], so that was really the only clue I could find.”*

In one case, the participant enumerated the devices they owned, but felt uncertainty around finding the offending device made the whole process meaningless. It is interesting to note that all participants experienced this kind of uncertainty, but only P14, who indicated they had technical expertise, felt it invalidated the remediation path: *“Can you see something useful, like an IP or MAC address or something? [...] I have no idea [what device could be the problem], so half of these steps I can’t execute. That makes this process kind of useless.”*

3.4.2. TAKING ACTION WITH A SUSPECT DEVICE

Only three participants reported that they were able to change the password of the suspect device (Fig 3.2). In these cases, the device either had an associated app or an interface on the device itself that allowed the user to initiate the password change. For, P11, who owned a Network Accessible Storage device (NAS): *“Yeah, resetting the password, you can do that via a small screen [...]. It worked, now with a slightly more difficult password.”*

Four other participants indicated that they thought the device did not have a password, e.g., P09: *“This [printer] has no password, does it? I can search on the internet, but I think the printer just appears on screen when I want to print. Other than that, there isn’t much to it. I don’t get any hits when I search for something related to passwords.”*

Four participants said they did not know how to change the password, as with P03: *“Well, I really have no idea how to do this. I do not have a booklet or anything. And the thing has no name, I think. So you tell me how to do this. A friend of mine helped me with installing this thing, but he got killed in a car accident, so I can’t ask him.”* One participant consulted the manual, P17: *“There is really nothing useful in the booklet that comes with it. I only see things that prevent us from suing them.”* Two participants reported visiting the manufacturer website, to no avail, as for P13: *“Yeah, I searched for this and I found a website that belongs to the device. But the site is totally unhelpful. I already know it is a camera, can’t they put something more useful on the site?”*

Two participants ‘solved’ the problem by completely disconnecting the device, e.g., P06: *“You know what, I will just disconnect it. I have no idea how to change the password, but it is broken anyway, so I will take it offline and then we will buy a better one [...] I don’t want a virus in my network.”*. P16 followed a similar behaviour: *“Well, I thought that [the camera] would hang there as a deterrent. But then I got your email. I threw out the device right away, because I definitely do not want a virus.”* Chalhoub & Flechais [60] considered disconnecting a smart device as a *compensatory behaviour* that owners apply to address security and privacy concerns, regardless of whether it directly addressed the concern.

When it comes to restarting the suspected device, two participants looked for a dedicated reset button. P10: *“I am pressing the reset button for a long time [...] OK, it is turning off and on again.”* The second person looked for such a button but ended up, like nine other participants, disconnecting the power cable: P02: *“I don’t really see a button or anything on the camera. Perhaps just pulling the plug then?”*

The last two steps concerned the modem/router. At least six participants had the standard router issued by the ISP. The email from the ISP contained a link to a help page that described two actions: how to restart the device by disconnecting the power, and how to factory reset the device via a web interface. While the email asked users to factory-reset the router, the presence of both actions on the help page led some participants to take the first listed action: only disconnecting the power. Strengthened by the presence of this action on the help page, participants were convinced their efforts were the requested ones, P02: *“It says here to pull the plug and wait for 10 seconds, I can do that, great”*. Moreover, participants tend to copy the actions they took for earlier steps and implement those for their router, P08: *“Reset? So I will do the same as with the camera. I have disconnected it for 5 seconds and it is back in. I see a green light so I guess that worked”*. Overall, 6 participants reported having enacted a factory reset, while 9 participants removed the power cable to reset the device.

P05, who was running a small business, said they did not want to execute a factory reset: *“The problem is that I would have to set up all port forwarding again and I don’t really want to do that [...] Then I have to let IT come again. [...] Were the previous steps not enough to make the virus disappear?”*

For the final step, 13 participants reported that they successfully set a new password via the ISP web interface of the device, while two said they did not know how to do this. For this step, six participants made use of the URL in the notification (see section B.1).

3.4.3. INFERRING THE SUCCESS OF REMEDIATION

When users manage to complete an action on the suspect device, they receive almost no feedback on the success of their efforts. The exception was when setting a new password was

supported via an interface that the participants are familiar with. The users who managed to reach a web interface for their router, for example, would get a clear confirmation when they successfully completed a password change. Still, all participants experienced actions that lacked feedback on whether they were successfully completed. More importantly, all participants lacked feedback on the success of their actions in terms of the main outcome: removal of the malware. These observations are of interest when compared to Forget et al. [126], and the examination of whether ‘engaged’ or ‘disengaged’ users arrive at secure outcomes to their (in)action to secure a computer – here the problem is that the outcome, secure or not, is not visible.

During the calls, we witnessed a clear desire by many participants to receive confirmation of whether they were doing the right things, as with P02: *“Shall I wait a few seconds? [...] OK, I think 10 seconds is enough, I am putting the plug back in [...] I am waiting for the lights to turn on again. It is supposed to be orange, right? Or green?”*

Some remediation actions were surrounded by uncertainty, while others were more clearly unsuccessful to the participants. In either case, participants regularly requested confirmation that they were successfully removing the virus. For instance, P04: *“Could it be enough if we do not change the password. That we do all other steps?”*, and P08: *“The device is already disconnected. Does that count as a reset if I now reconnect it again? I am really curious whether the virus is really gone. Can I reconnect it now?”*.

3.4.4. MOTIVATION UNDER UNCERTAINTY

All participants were willing, in some cases eager, to undertake the recommended actions, e.g., P09: *“I am now putting the plug of the router back in. What is the next step of this adventure?”* Participant motivation was illustrated by the degree to which they tolerated their uncertainty about what was asked of them, and whether they conducted the actions correctly. Motivation was also visible through the effort that was made. For example, the device or router might be in another part of the house or access to it might be blocked. This was the case for P03: *“You ask quite a bit from me, because then I have to make quite a mess. [...] Let me put the phone down, I need to move a few boxes... OK. What do I do now?”*, and P07: *“Then I will walk to the utility closet [...] I see the cable already, I will pull it out completely.”*

In addition, the factory reset of the router means that users lose their configuration, which might not be trivial to set up again. P10 debated this, *“Ah, so then I have to set up all port-forwarding and port assignments again. Well, I think that is the right thing to do, otherwise the virus will hang around.”*, as did P04: *“Oh, that is complicated. I did the same thing a while ago, but then I need to reconfigure all port forwarding again. But OK, if that*

helps, then we will do it again.”

Only a few participants expressed doubts about the effort, in all cases because they were not clear what problem Mirai posed, as with P01: *“Eh, let’s take a step back. I have no idea whatsoever about how that Mirai virus actually works. I mean, I do not experience any issues, right? So what is the problem?”* After an explanation about how Mirai-infected devices are used for criminal activity against other users and organizations on the internet, P01 concluded: *“Ah, right. That is understandable, I am happy to cooperate.”*. Renaud & Goucher [238] note that the ‘gulf of evaluation’ differentiates between the sense of being able to enact a security behaviour, and the ‘response efficacy’ of whether the behaviour is appropriate.

No participants dropped out before completing the steps. The only case where a participant did not want to conduct a specific step was P14, who felt none of their devices were plausible suspects, and as such did not want to implement a reset and password change on any of those devices. They did, however, proceed with subsequent steps involving the router.

Regarding the evidence for users’ seemingly high motivation, one potential source of bias here (as discussed in Section 3.3.7) could be an observer effect (a.k.a. the ‘Hawthorne effect’), where the fact that the participants know their actions are being ‘observed’ makes them more motivated than they might have been without the presence of the researcher.

3.4.5. THE END: REMEDIATION, AND REINFECTIONS

Table 3.1 presents an overview of participant-level actions and outcomes. Again, the coding used in the columns for the remediation steps relate to the boxes in Figure 3.2. After the intervention, 14 of the 17 participants were observed to be remediated, as measured by the absence of their IP address in the daily data feed of Mirai infections received by the ISP in the four days after the call. This may count as good news. The cumbersome non-deterministic remediation process seems at least probabilistically related to the desired outcome. Three participants remained infected. It is true that they did not fully execute the recommended steps, but the same holds for other participants who were regarded as having managed to remediate. Only four participants could be said to have fully executed the recommended actions (P01, P10, P11, P15). We include P15, because this person took the suspect device permanently offline, so in that sense ‘secured’ it from further harm. We monitored the presence of the IP address in the daily data feed for two more weeks after the remediation period. In 5 cases, we observed a re-infection with Mirai; there was a gap of three consecutive days where the user’s IP address was not reported in the daily data feed, and then it reappeared. Two of these reinfections were non-remediated users, three were users who did manage to remediate at first.

For the two non-remediated cases, the infection disappeared by an unknown cause five or more days after the call. This is consistent with the relatively high ‘natural’ cleanup rate seen elsewhere [55]. One explanation is that the Mirai malware is not persistent on the device, at least not at the time of the study. This means that a power cycle may have removed the infection, although the device is still in a vulnerable state. It might be discovered and reinfected soon thereafter, because of the aggressive scanning conducted by Mirai bots.

The three cases where we observed an initial remediation, and a later reinfection, can have various explanations, and as such are indicative of avenues for future work. One explanation is that the detection of infections via the daily data feed is not perfect, potentially including false negatives. Another explanation is that these users did manage to get rid of the infections by power cycling the devices, but did not remediate the underlying vulnerability (i.e., set a secure password). This is consistent with our observations, because all three users did not fully execute the recommended actions. As noted from the observations, users may have otherwise had multiple infected devices and only focused on one, or focused their attention on the wrong device.

In the end, the gap we observed between advice and user actions cannot be blamed wholly on either the user or the advice-giver the ISP. It points to the responsibilities of a third actor: the manufacturer. Even when users went online and tried to find manufacturer information about solving security problems, there were complications. This was certainly the case for P16, who was not able to even identify the manufacturer: *“Well, there is no brand name on the device, haha, only IP-camera is printed on the side of it.”*

3.5. DISCUSSION

Returning to our overarching research question, we provide real-world evidence of the gap between advice and outcomes in IoT [34], but also the impact this gap can have on smart home users. There are two sides to this story – the quality of advice, and the characteristics of the response to that advice.

Successful behaviour for our participants was often unconfirmed and unconfirmable, and neither the users nor the advice-giver can resolve this at present, given the constraints inherent in the situation (in home infection, limited device visibility, etc). This unbridgable gap points to the responsibilities of other actors, notably the manufacturer [140]. We could argue users lack capability, but it is not a lack of user capability, but a design flaw, pointing to the relationship between behaviour support and interface design to provide situational feedback (as highlighted elsewhere for user access control guidance [304]). The lack of ‘normal’ computing interfaces on IoT devices creates an environment fraught with confusion and uncertainty for applying standard security advice.

What we have for network-connected smart home devices is also a multi-party intervention. Participants had to wait for their efforts to be confirmed as worth it (that remediation will be confirmed at some point afterwards via network scans, and a *lack* of capacity for the ISP to follow up). Participants demonstrated despair over not knowing what to do and whether their effort was successful. Remediation is then non-deterministic (very likely to work, but not definitely going to work). The lack of feedback stands in contrast to, say, removing Windows malware, where a removal tool—such as an anti-virus client—will typically report on what it found and whether it was effective in removing it. This limits the potential of checklists, for instance, if instructions cannot be made specific enough to a particular user's set of network-connected devices (and are as such, 'sub-optimally targeted').

Participants applied one of the heuristics identified in our results, to navigate the gap in specificity, and attempt to identify the target of an advocated behaviour. Applying advice then leans on motivation, in that most participants were willing to try quite convoluted steps (going to another room to unplug the router, coming back to the phone, then back to the router, etc.). Where Redmiles et al. [232] isolate 'bad advice', we step back from this to identify 'ecosystem factors which limit the capacity to construct good advice'. We regard this then as also exploring the limitations of *emergent* interventions for smart home security.

What is remarkable and worthy of further exploration is that our participants demonstrated somewhat correct reasoning in identifying suspect devices, consistent with actual properties of these devices. Mostly the heuristic is to eliminate suspect devices. This further highlights the importance of local context to instantiating security advice for the smart home [304], but also making advice specific enough to be actionable [237].

3.5.1. INFORMING EFFECTIVE INTERVENTIONS

Where participants felt a need to enumerate over familiar behaviours, many would push back if they did not know how to enact the advocated behaviour. This points to self-efficacy, important for prompting action within various behaviour change approaches [112]. To put our findings in the context of enacting (what appeared to our participants as) a new behaviour, acting on notification of a malware infection is an *opportune moment* or *prompt* to enact a new behaviour, so we refer to the Fogg Behaviour Model (or B=MAP / B=MAT model [123]). In this model, Behaviour = Motivation + Ability + Prompt. The model has been used extensively across areas such as persuasive design and personal development, but also to understand social interventions for security [77], and opportunities for security interventions in a retail environment [217].

A Prompt can be a Facilitator, Spark, or Signal – here it is a Signal, that a device

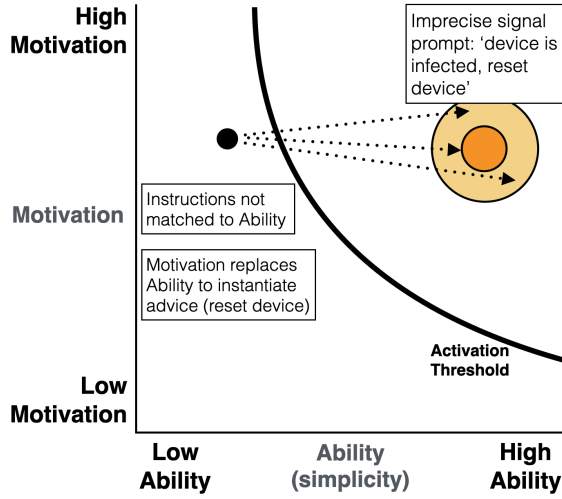


Figure 3.3: Action Diffraction for resetting a smart home device. Users may vary in Motivation, and rely on their Motivation to enumerate over possible solutions (standing in for a lack of knowing the precise Ability they need to apply). The target behaviour may be deterministic (the small circle, top right), but plausible variations surround it, informed in part by Instructions. It can be unclear if the applied Ability has achieved the intentions of the Prompt, even if it has been successful.

in the home is infected and that actions must be taken to resolve the issue, as a call to Motivation and upon an Ability to act. The ISP carries the Signal to the user (highlighting that ISPs are the *best-placed* party to intervene, but that this does not mean they are the *most appropriate*) – this relies on sufficient Motivation and Ability already being present. We found that participants were *over-investing* Motivation to make up for an insufficient definition of the target behaviour or outcome (a lack of capability to identify or confirm the appropriate Ability). Among our participants, there was uncertainty as to what was right to do, to the extent that a user may enact a behaviour which removes malware, but continue with further actions for lack of indication that they had already succeeded. This even includes where some of our participants chose to permanently disconnect or dispose of a suspect device (representing an *unintended harm* of unclear advice [63]). ‘Actionable choices’, with clear outcomes, are regarded as feasible in areas such as smart home privacy [251], and in supporting a user-defined ‘recovery state’ [143].

We show the gulf where these harms manifest as what we refer to here as ‘Action Diffraction’ (Figure 3.3). Where Renaud & Goucher refer to the ‘Gulf of Execution’ [238] (including knowing what needs to be done, but not how to do it), here we find a gulf created by restrictions in the vehicle of the intervention itself which makes the target behaviour indistinct. This applies to both knowing what the target behaviour is, and knowing whether it

has been reached. Where the Activation Threshold is the point of realising a target Behaviour, and a user being activated to try to get over the Threshold, our results show efforts being ‘diffracted’, splitting off in many directions as participants find themselves exploring non-deterministic and potentially inapplicable behaviours (this includes where they have Ability to do something, but are not willing to try everything unless they can be Motivated to do so).

Renaud & Goucher [238] frame a ‘Gulf of Evaluation’ in formulating an intention to adopt a secure behaviour, and Redmiles et al. [234] identify dimensions of advice quality. We note in reference to the latter that the specificity – and actionability – of advice, including the capacity to evaluate the efficacy of the behaviour [238], are also impacted by the specificity of the target behaviour and its confirmation. Our findings showed also that, as with other forms of security advice [237], multiple sources of instructions can potentially confuse users further.

One contributory factor to this problem is best articulated through the Behaviour Wizard of Fogg & Hreha [125]. The best-practice advice seen by smart home users is an ‘unfamiliar’ task (requiring a link to existing practices), but framed more like a ‘familiar’ task (one that does not need explanation), and so we saw participants replacing an unfamiliar action with familiar behaviour(s). This is a complex world of Things, where enacting the wrong behaviour can result in ‘proxy changes’ [214], regardless of whether the intended outcome is reached. A user may turn on and off many devices, or the wrong one and not the right one, or achieve the goal but lose tailored configuration settings in the effort, all while not knowing in the moment whether they have succeeded.

3.5.2. IMPLICATIONS FOR EVOLVING IOT THREATS

If users only apply some of the advice they are given, or devices have inherent security weaknesses, they may *continue* to be vulnerable and require *regular* intervention. Users can follow advice but still suffer the same consequences again, if IoT infrastructure does not help them to stay recovered, or malware evolves. There are parallels to the Transtheoretical Model [224], where understanding specific stages of behaviour can identify security improvements [221]. Inherent weaknesses in the design of many smart home devices put a user back into an ‘unhealthy’ situation (e.g., a device repeatedly falling back into an insecure state), requiring repeated cycles of *contemplating* and *acting* on advice, to *maintain* secure devices.

New malware variants are moving away from short-lived infections, and becoming persistent and resistant to current interventions [47]. More efforts of the type we have observed for Mirai infections would be required where, for instance, thousands of QNAP network access storage devices have been targeted by persistent malware [303], and the direction of travel shows that advice from manufacturers is requiring users to follow 20 or more

steps completely and successfully to resolve these issues [228]. Moreover, some of these variants are also starting to include countermeasures to make detection difficult. For instance, malware leveraging blockchain DNS or TOR makes it even harder for the interveners to assess the efficacy of the user's actions [40, 276, 287]. This is all within the context of increasing use of smart home devices, which itself already increases the complexity of remediation when there are problems (as we saw evidence of here).

3

3.5.3. RECOMMENDATIONS

Here we describe recommendations emerging from our Results and consideration of behaviour change approaches, associating recommendations to specific stakeholders.

- **Confirmation of settings changes.** Visibility of changes to system status is a crucial design principle [207]. Here this applies to both Apps and Interfaces, as created and maintained by the *manufacturer*. This was seen among our participants as already happening for some devices and interfaces, but should be enshrined as a consistent design choice, to reduce the 'diffraction' of remediation efforts. This would then serve as a visible 'security outcome' [126], to *then* be able to consider whether the visible outcome was the correct step to follow. This may be necessary for future security issues if resetting / unplugging a device actually runs the risk of reinstating default credentials, for instance. This would complement efforts to standardise smart home device functionality (as in e.g., the UK [277] and US [117]) which aim to have *manufacturers* reduce the scope for misconfiguration as a vector for device compromise (as with e.g., easily-guessed 'default' settings).
- **Settings logs.** A log of settings changes can help both *users* and *ISPs* (or indeed anyone 'helping' users) to see and refer to a clear record of changes. This could also include notifying users of security settings which need to be changed at setup but have not been, or which have been changed but not by a registered user. Ideally, there would be some signalling to users when a security issue is suspected, where there is a general lack of event logging related to security [117].
- **Assisted remediation.** Our study showed that not all participants were able to follow the advice, or needed confirmation that they had followed it. For lack of being able to move incrementally toward a clearly focused outcome (Figure 3.3), having a helper on-call or on-site would increase chances of a successful outcome, if the previous steps cannot be achieved. This would be a low-bar in terms of ensuring that there is an intervention for all levels of Motivation and Ability – if *users* are as keen to follow

advice as our participants, they cannot be blamed if they are trapped in a cycle of trying advice without confirmation of actions or visible evidence of success. This relates to having actionable choices to begin with. It also aligns with the incentives of ISPs, which could commercially offer such services, though this brings the risk of users distrusting notifications as a ploy to sell a service – ISPs might only offer the service if the user asks for it.

3.6. RELATED WORK

Chalhoub & Flechais [60] studied real-world users of smart home devices, where limitations in device features and transparency were seen to frustrate privacy-related decisions. The authors characterised *compensatory behaviours* in response to concerns (such as disconnecting a device). We saw participants defaulting to ‘familiar’ behaviours as a strategy to approach the uncertain process of situating generic advice. Geeng & Roesner [133] studied multi-user smart homes, noting that when devices fail to function properly, alternative paths to using a device are needed. We saw a parallel, where participants sought a viable solution to critical security issues, but were at times reluctant to dismantle their smart home device configurations to achieve it.

In terms of supporting behaviour change, Forget et al. [126] studied the security attitudes, behaviours, and understanding of active computer users from device activity and interviews. The authors characterised ineffectively proactive users, who exerted too much effort for security or regularly performed familiar behaviours even if they did not match the security concern. Where the authors saw information-seeking behaviours, our participants felt challenged in determining what to seek information about (lacking both clarity as to what was the target device, and available diagnostic information). Crucially, Forget et al. highlighted the importance of tangible outcomes to user actions, where here there was a lack of clear outcomes; the authors identified ‘problematic knowledge gaps’, where for consumer IoT environments these gaps are constraints in advice and user support.

Reeder et al. [237] identify a range of criteria for good home security advice, including that it must be actionable. We identify a gap that requires the recipient of smart home security advice to be able to complete advice and relate advice received from others to their personal context. The authors also discuss the potential need to enumerate over possible versions of generic advice to reach specific advice, considering “offering the generic advice followed by specific instructions on how to implement it” – similarly, our participants applied strategies to do this themselves.

Redmiles et al. [234] identify ‘perceived efficacy’ of advice as important, where here there is an element of efficacy in being able to localise advice received from others. The

advice the authors reviewed was regarded as mostly ‘actionable’, where here we explore the implications of advice which, at least for our participants, was not immediately actionable. Redmiles et al. regard network security as amongst the least actionable and most general security advice (e.g., “Secure your router”), raising questions of whether non-actionable advice should be given to users in the first place, and we provide real-world evidence informing this discussion.

Cetin et al. [54] studied a ‘walled garden’ approach of limiting users’ capacity to access the Internet while a device is infected. Here we learned about the remediation journey while users were acting on suggested remediation actions locally themselves, rather than checking the effectiveness of the notification method alone.

3.7. CONCLUSION

Here we studied user efforts to apply advice provided to them by their ISP. We found that the advice was not specific enough to ensure that it was applicable to participants’ own smart home context. Critically, constraints to the specificity of advice limited how it was produced, communicated, and put into practice in a real-world setting. Only 4 of 17 participants completed all applicable advice steps successfully. Action typically went wrong at the second step (changing the password of the suspected device), or at the fourth step (resetting the router to its factory settings). 16 participants were able to pinpoint a plausible infected device, using one of three strategies we identified (including by process of elimination).

Our work informs the understanding of interventions for real-world IoT settings. The construction, communication, and enactment of technical advice to home users is both complex and collaborative. It involves end-users, their ISPs, device manufacturers, and technical experts to support successful outcomes. Putting our findings into perspective with the continuing need for technical support for home computers and mobile devices, the need to fix security issues of smart home devices can be expected to persist. Given the complexity and role of local context, this can be expected to require analysis of the smart home in situ, including return visits to users of reinfected devices. Future work will explore the capacity of intervention approaches which include multiple relevant stakeholders. For instance, a list of known vulnerable device models could aid both ISPs in informing end-users, and end-users themselves in identifying problematic devices which they use or are considering for purchase.

4

REMEDIATING PERSISTENT IOT MALWARE

Consumer IoT devices may suffer malware attacks, and be recruited into botnets or worse. There is evidence that generic advice to device owners to address IoT malware can be successful, but this does not account for emerging forms of persistent IoT malware. Less is known about persistent malware, which resides on persistent storage, requiring targeted manual effort to remove it. This paper presents a field study on the removal of persistent IoT malware by consumers. We partnered with an ISP to contrast remediation times of 760 customers across three malware categories: Windows malware, non-persistent IoT malware, and persistent IoT malware. We also contacted ISP customers identified as having persistent IoT malware on their network-attached storage devices, specifically QSnatch. We found that persistent IoT malware exhibits a mean infection duration many times higher than Windows or Mirai malware; QSnatch has a survival probability of 30% after 180 days, whereby most if not all other observed malware types have been removed. For interviewed device users, QSnatch infections lasted longer, so are apparently more difficult to get rid of, yet participants did not report experiencing difficulty in following notification instructions. We see two factors driving this paradoxical finding: First, most users reported having high technical competency. Also, we found evidence of planning behavior for these tasks and the need for multiple notifications. Our findings demonstrate the critical nature of interventions from outside for persistent malware.

4.1. INTRODUCTION

Smart home devices keep proliferating and, unfortunately, so do the malware families targeting these devices. Solutions for malware detection and removal have a long lineage, going back at least two decades. After the chaos of the global virus outbreaks of the early 2000s, countermeasures slowly started to emerge from what became the anti-virus industry and from operating system manufacturers like Microsoft. Years of painstaking development have resulted into the highly automated and usable tools that consumers rely on today to detect and remediate infections on their personal computers.

Then, about five years ago, IoT malware surged, most notably in the form of the Mirai family [22]. It captured millions of surveillance cameras, digital video recorders, routers, and many more devices that researchers could not identify [243]. Here, none of our automated tools work and many of the hard-earned usability lessons for PC malware cannot be applied. These devices are typically headless, lack a graphical user interface (GUI) or a peripheral device for users to learn about an infection and take the recommended actions. There are no standard anti-virus tools that can run on these devices. (As an aside, some vendors are now offering dedicated anti-IoT malware devices which users are meant to put in their local network. Bundled with a subscription, they can cost hundreds of dollars per year, which explains why they currently are niche products. These devices can potentially do detection based on network traffic, but not remediation of the infection. The latter task remains with the user.) To make matters worse, IoT represents an enormously heterogeneous population of devices in terms of design and function [166]. The deployment of tens of thousands of different devices makes it all but impossible to give users security advice that is actionable for their specific devices.

Industry and governments have been coping with this challenge by providing consumers with highly generic instructions that try to cover all manner of devices and attack vectors [196, 198]. This advice suffers from usability problems, since it could not be made actionable for specific device, leaving consumers to figure out how to take actions like disabling Telnet, changing a factory-default password or installing a firmware update. Surprisingly, these coping strategies did have some success.

Remediation levels were found to be high [55], even though the security advice was poorly understood by users and it lacked a deterministic path to removing the infections [41]. This success was helped by fleeting nature of the infections. The bulk of all IoT malware resides in memory only and does not gain a persistent foothold on the device. Thus, a power cycle would remove it—albeit only temporarily if not combined with other protection measures like a password change.

Now the next challenge has arrived: persistent IoT malware [40, 47, 276, 287]. It

combines the worst of both worlds: the persistence of PC malware with absence of effective and usable detection and removal tools of IoT malware. Does persistent IoT malware make remediation more difficult? How do users experience their remediation efforts? Learning the answers to these questions is critical in responding to the next evolution of IoT malware.

This paper presents the first field study on removing persistent IoT malware by consumers. We partnered with an Internet Service Provider (ISP) in The Netherlands to compare the remediation times of 760 customers for three categories of malware families: persistent Windows malware, non-persistent IoT malware and persistent IoT malware. In the latter family, we focus on QSnatch, also known as Derek [196], as a case study. We selected QSnatch since it was the most prevalent IoT malware family, which was not memory resident only, in the network of our partner ISP at the time of this study. Besides, QSnatch is an appropriate representation of a persistent IoT threat for several reasons. First, according to the US Cybersecurity Infrastructure Security Agency (CISA) and the National Cyber Security Centre UK (NCSC-UK), the number of QSnatch reported infections grew from 7,000 devices in October 2019 to more than 62,000 in June 2020 [64, 76]. Second, non-profit organization Shadowserver recently reported QSnatch as the second most important threat after Mirai—in some countries even as the top IoT malware family [254]. Finally, as highlighted by [17] network attached storage (NAS) are among the top devices targeted by IoT malware.

Next, we contacted ISP customers who had suffered from a QSnatch infection in the past year. We interviewed 57 customers with an instrument design informed by the COM-B behavior model [185], which stresses the importance of individuals' capabilities, motivations, and opportunities to perform a behavior. The model has been suggested to be applied to understand behavior change in security [103]. We also compared the cleanup success of interviewees to the non-interviewed QSnatch victims.

Overall, we find that, yes, persistent IoT malware is more difficult to remediate. These infections last more than three times longer than infections with Windows malware or non-persistent IoT malware, namely Mirai. This is consistent with the fact that the remediation for QSnatch consists of a convoluted series of steps. Surprisingly, though, the interviewed users reported that they did not find the remediation steps particularly difficult.

We see two factors driving this paradoxical finding. First, most users reported having high technical competency—in fact, the majority even reported working as an IT professional. So their tolerance for difficult tasks is a lot higher than for the average user. Their frame of reference might be complex IT admin tasks, rather than the more simple consumer action of running an AV tool. We found evidence of planning behavior for these tasks, e.g., doing it on the weekend. There might be a self-selection process at work, owners of network-attached

storage (NAS), as a new technology, are much more likely to be technically competent [84], thus experiencing the difficult task as a normal task, but then they do need some time and effort to execute it.

The second factor that explains why users did not find it very difficult, yet they took longer to remediate than Windows and Mirai infections, is that the latter can also get remediated without user action. An automatic scan of an AV tool or Windows malware removal might find and remove the infection, without the user even noticing. For Mirai infections, a power cycle removes the infection (even though it leaves open the possibility of reinfection). Such ‘natural’ remediation is not possible for QSnatch. Only user action can get rid of it.

In sum, we make the following main contributions:

- We quantify the infection duration for 228 customers infected with persistent IoT malware, and compared it to customers infected with memory-resident IoT malware and Windows malware. Compared to Windows malware, the mean infection time of persistent IoT malware is three times higher. Compared to memory-resident IoT malware, persistent IoT malware mean infection time is five times higher.
- We estimate the survival probability of different types of malware and statistically compute differences between malware families. Our results show that 30% of the infected subscribers with persistent malware remain infected even after six months since the infection was detected.
- We provide real-world evidence of users mitigating persistent IoT malware. Our results show that all QSnatch-infected customers remediate right or closely after receiving a notification.
- We derive a set of recommendations to expedite the cleanup processes of persistent IoT malware based on the interviews conducted with 57 infected customers.

4.2. BACKGROUND

4.2.1. QSNATCH AND PERSISTENT MALWARE

Most popular IoT malware families, such as Mirai in its many variants, are stored within the temporary file systems of the IoT devices. They resided in the Random-Access Memory (RAM) of devices. This memory is volatile, thus any malware residing in it will be removed from the device if the device is powered off or just restarted. Persistent malware, on the other hand, is stored among system files of the operating system, they can be added to the

startup process of the operating system or schedule processes, being able to survive reboots, and maintaining a connection with the device to keep it as part of a network of bots.

Our study focuses on an important example of persistent IoT malware called QSnatch. QSnatch targets network-attached storage (NAS) devices from the manufacturer QNAP [226]. Some characteristics that make QSnatch persistent are that it changes scheduled tasks of the device, prevents firmware updates by rewriting the URL from where the update comes from, and steals usernames and passwords [79]. The malware uses Domain Generation Algorithms (DGA) to communicate with the command and control servers controlled by the attackers [51].

Different security firms have characterized the capabilities of QSnatch [64, 262]:

- Common gateway interface (CGI) password logger - a fraudulent version of the device admin login page, recording authentications and passing them to the legitimate login page.
- Credential scraper.
- SSH backdoor - Allowing to execute arbitrary code on a device.
- Exfiltration - When run, QSnatch steals a predetermined list of files, which includes system configurations and log files. These are encrypted with the attacker's public key and sent to their infrastructure over HTTPS.
- Webshell functionality for remote access.

QSnatch poses a threat to users besides the possibility of being used for Distributed Denial of Service (DDoS) attacks or to deliver malware payloads to other devices.

4.2.2. QSNATCH REMEDIATION MECHANISMS

To remediate QSnatch, QNAP has published a series of recommended steps. Our partner Internet Service Provider (ISP), in turn, created a shorter version of these steps to include in their notifications to affected customers.

MANUFACTURER RECOMMENDATION

QNAP's security advisory to address QSnatch infections recommends a whopping 84 user actions, organized around several high-level steps [227]:

- Update QNAP turbo station (QTS) to the latest available version.
- Install and update Malware Remover to the latest version.

- Install and update Security Counselor to the latest version.
- Update your installed QTS applications to the latest versions if available in the App Center.
- Configure settings to enhance system security.

Each of these steps includes actions like changing various settings of the device, enabling and disabling features, changing passwords and configurations, and subscribing to QNAP Security Newsletters [227].

Different than how Mirai could, in practice, be removed by resetting an infected device [41, 55, 244], resetting a NAS would not lead to remediation. In contrast to Windows malware, where users may count on existing tools to remove infections in the background, such as antivirus software, removing QSnatch requires recognizing the correct information and applying the security advice. To solve the issue, users need to perform more steps than for removing Mirai malware [41, 55, 244] or running antivirus, and if these tasks are perceived as challenging or dull, users might postpone them [279]; thus, making this infection more difficult to remediate.

INTERNET SERVICE PROVIDER RECOMMENDATION

The partner ISP contacts customers who suffer from a QSnatch infection. The notification includes the recommended steps for remediation. Rather than point customers to the complicated advisory on the QNAP website, the ISP has condensed the advisory into a shorter and simplified version of the remediation process.

The notification explains to the user that a QNAP network-attached storage device has been compromised with QSnatch malware and then provides nine steps to solve the infection (see C.1 for the full notification). Since QSnatch has the capability of rewriting the URL for downloading the new firmware and blocking the launch of the QNAP Malware Remover tool[79], the ISP recommends to users to do the following:

- Go to the website: qnap.com/en-en/download
- Under “1 - Product type”, select the option “NAS / Expansion” and select the number of slots present on the right.
- Under “3-Model”, select the type of NAS you are using.
- Under the “Operating System” tab, select the most recent version and download it via the “[REGION]” button.

- Open the NAS on your PC or Mac and choose firmware update, and then Manual update.
- Browse to the downloaded file and update the firmware / operating system.
- Go to APP Center and choose “Malware Remover” and download it on your PC or Mac.
- Click on “manual update” in App center, browse to the download file and update the Malware Remover.
- Run a scan with the Malware remover.

NOTIFICATION PROCESS

At our partner Internet Service Provider (ISP), the starting point to handle all infections is a feed from a third party, Shadowserver, specifically the Drone Report [253]. Shadowserver is a non-profit security organization that shares abuse data to make the Internet secure. It is a trusted source for network providers, national governments and law enforcement [7]. The Drone Report is received daily by the ISP abuse handling department and contains data on infections for many different malware families.

It includes the Internet Protocol (IP) addresses where infected machines were observed. These addresses were captured by different techniques such as sinkholes, darknets, honeypots and other sources [253]. The ISP connects the Shadowserver IP addresses with their customers’ data. Once the affected customers are identified, an automated system sends an email notification about the detected security issue. These notifications can be customized to the type of malware. So for Windows malware, users get different instructions than for Mirai IoT malware.

Shortly after QSnatch was added to the Shadowserver Drone Report, the ISP included these infections in the standard abuse handling workflow for all infections. Affected users would receive the email with the ISP’s customized recommendations for removing QSnatch.

4.3. METHODOLOGY

Our study was built upon the existing process of our partner ISP, which includes identifying QSnatch infected users and notifying them as shown in Figure 4.1. Our mixed-method approach started from the survival analysis of 760 customers for three mutually exclusive categories of malware families: Windows malware, non-persistent IoT malware and persistent IoT malware. These infections were identified and tracked using the daily reports

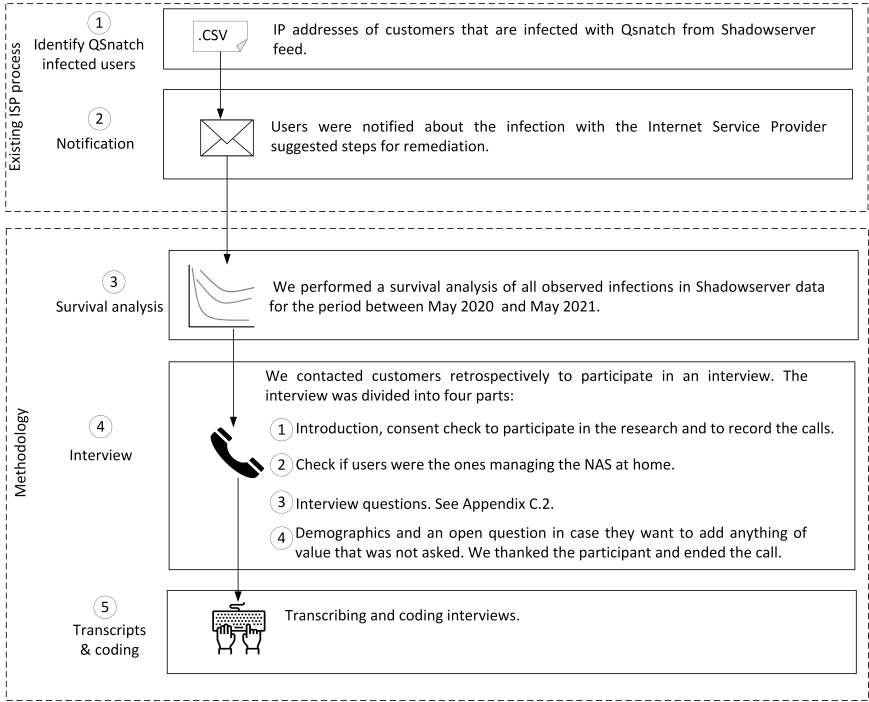


Figure 4.1: The ISP existing process and overview of the mixed-method approach

provided by Shadowserver [253] which were recorded in the abuse department system of the ISP from the period between May 2020 and May 2021.

Thereafter, we contacted customers infected with QSnatch malware to carry out an interview to understand how they handled the infection. The interview design was informed by the COM-B behavior model [185], which stresses the importance of individuals’ capabilities, motivations, and opportunities to perform a behavior. The model has been suggested to be applied to understand behavior change in security [103]. Finally, the recordings of the interviews were transcribed and coded for its analysis.

4.3.1. SURVIVAL ANALYSIS

DETERMINING INFECTIONS

As described in the section 4.2, our partner ISP connects the IP addresses flagged in the Shadowsever “Botnet Drone Report” [253] with their own customer data to determine which customers are infected. The Drone Report contains data on infections for 112 different malware families, including QSnatch, Mirai and Windows malware. Shadowserver reports are generated daily, and there were no gaps between reports during the period of this study for our partner ISP.

Once infected customers have been identified, a case is created in their incident ticketing system to follow up with customers, and send notifications. A case is closed once the IP addresses belonging to an infected customer are not seen again in subsequent reports. Note that a customer might receive multiple notifications for the same infection if the infection persists for long periods of time.

In this research, we track users' infections with customers IDs. These IDs are unique and can be associated with multiple IP addresses over time. This way we can track the whole period of infection, even though the IP address of the customer might change in the course of the measurement period.

In the normal workflow, the notifications are sent the day after the Drone Report has reported the infection. However, during the period of this research, the abuse department was transitioning to a new system. This caused delays in sending out some of the notifications. These were randomly distributed across the infections in the Drone Report, thus across all types of malware. The transitioning of the abuse system does not affect the results of this research; instead, this served as a natural experiment. A natural experiment is where a circumstance that was not controlled by the researchers occurs giving the opportunity to evaluate the intervention [168], in this case, the impact of notifications. All infected customers were eventually notified, either the day after the IP address was first detected by Shadowserver or after a second or third detection.

Our starting point is a set of 760 customers that got notified for three categories of malware infections: Windows malware, non-persistent IoT malware (Mirai) and persistent IoT malware (QSnatch). These notifications were sent over the course of a year: May 2020 to May 2021. The dataset includes 228 customers infected with QSnatch, 107 customers infected with Mirai, and 425 customers infected with Windows malware. For Windows, infections consisted of the following malware families: Ramnit, Kovter, Citadel, Qrypterrat, Conficker, Necurs, Sality, Caphaw, Downadup, Emotet, Gamarue, Gozi, Necurs, Nivdort, Nymaim, Grypter.rat, Ramnit, Sirefef, Tinba, and Zeroaccess. These Windows malware families have a wide range of capabilities, from banking trojans to worms to ransomware. We group all the Windows malware families together, rather than comparing individual families. First of all, we are interested in comparing malware categories—persistent/non-persistent and IoT/non-IoT—rather than individual families. Second, for some families the sample size was very small (e.g., 1, 4, 7 or 8 observations). This rules out meaningful comparisons.

KAPLAN-MEIER ESTIMATES

We selected all infections with a start date between May 2020 and May 2021. To unravel how cleanup of QSnatch infections compares to other malware categories, we computed the survival time probability for customers infected with QSnatch, Mirai, and Window malware

families. For this purpose, we used Kaplan-Meier (K-M) curves and estimates [157, 239]. To construct the survival time probability and curves of the different malware families, we used the starting date of the infection and the final date of infection from the historical Shadowserver data stored in the incident ticketing system of the ISP (see ‘Infection Time’ for more details). As described by [239], this allowed us to compare all the observations within the groups and begin the analysis at the same point, we check their lifetime until cleanup occurs or the observation period ends. The latter cases are censored. Censoring means the total survival time for the observation cannot be precisely determined since it falls outside the period of data collection [239]. These data points are retained in the analysis, but they are considered as the event did not happen. In this research, observations identified during the last 14 days of the period of observations were right-censored.

The Kaplan-Meier estimator is a non-parametric statistic used to estimate the survival function. The function is defined as: $S(t) = P(X > t)$ [160]. In this equation, S is the probability that a random variable X , in this case that malware is still on the device, exceeds a specified time t . We used the lifelines library [50] to plot the curves and compare them visually and statistically.

To statistically test whether the differences between survival curves are significant, we use the log-rank test [36]. This is a method to compare the survival functions of different populations. It compares the estimates of the hazard functions of two groups at each observed event time.

INFECTION TIME

To construct the survival probability, we needed to estimate the duration of infection. We used a year of historical infection data that the ISP receives from Shadowserver [253], so we consider as “infection time” the period between the first time the infection is detected until the last time the infection is seen. For all cases, we had the starting point of an observation, but in some cases, we could not determine if the remediation happened or not since the infection was detected close to the end of the period of observation. Thus, observations identified during the last 14 days of the period of observations were right-censored.

Note that in the survival analysis of the interviewed users, only 55 observations will be presented. Due to the system transitioning of the abuse department for two customers, we could not retrieve the closing date of the infection. In other words, these two users were notified since they were seen in Shadowserver and added to the incident ticketing system; thus, we contacted them for the interview, but the ISP system did not record the end-time of the infection to calculate the survival probability.

Note that the infection time as observed in Shadowserver consists of the time it took the ISP to notify the infected customer, the time the customer waited before taking action, the

time it took to execute those actions, and the time the infection remained on the device if the actions were unsuccessful in removing the infection.

To avoid any confusion, we should note that during the interviews we asked users if they could roughly estimate how much time they took to perform the remediation steps. We consider this time the users' self-reported time of dealing with the infection. This should not be confused with the total time of the infection, as derived from Shadowserver observations.

4.3.2. INTERVIEWS

To understand the process that users follow to perform the steps, and determine if they are able to deal with persistent malware, we developed an interview protocol which we executed in April and May 2021, at the end of the observation period for infections. The interview was a structured interview with closed questions with the opportunity to elaborate on the answer, and some open questions.

The downside of contacting customers retrospectively is that there was a time difference between the interview and users' actions, which we will discuss more in subsection 4.6.4 (Limitations). Also, we chose for Qsnatch-only interviews, rather than a design that would have interviewed people from all three "treatments" (Qsnatch, Mirai, Windows). This choice has pros and cons. We acquired more data on the challenges of a new and non-studied group, but we cannot compare the answers of the different groups and connect them to the different remediation speeds.

From the total set of customers who suffered a QSnatch infection in the year May 2020–May 2021 ($n=228$), 45 (20%) were contacted to carry a pilot to test the protocol (See Table 4.3.2). Then the remaining customers, 183 (80%), were invited to participate in an interview via email. The email stated to customers that they were notified in the past about a QSnatch infection and that we wanted to learn about the actions they took, if any, to remediate the infection. Of the 183 customers, 57 (31%) accepted to participate in the survey. We later checked for selection bias by comparing the remediation rates for the interviewed users versus the non-interviewed users and found no significant difference.

The interview was divided into four parts as described in Figure 4.1. First, we obtained consent from the users to participate in the study as well as recording the interviews, and users were reminded that they could step out at any time. Second, we asked if the person we contacted was the one who manages the device, if they received the notification and if they understood the notification. Third, different questions about how users handle the infection were asked. This design was informed by the COM-B behavior change model [185]. More details can be found in this section, regarding how the principles of COM-B were seen as useful to the study, and how the questions within the interview protocol were based on the

COM-B pillars.

Finally, a number of demographic questions were asked, as well as a closing question in which users could add any remark that was not covered during the interview. Next, we thanked the participant for his time, and finished the interview. See the complete interview protocol in C.2. The recordings of the interviews were transcribed and coded using ATLAS.ti software.

COM-B AND SECURITY BEHAVIORS

The COM-B model has been proposed as applicable in the goal of understanding motivators and blockers for secure user behaviors, both for home users and in organizations [103]. The pillars of the model (Capability, Opportunity, and Motivation) act as attributes which must all be in place to provide the conditions for a behavior change intervention to be regarded as complete. As we describe in subsection 4.2.2 QSnatch cleanup is complex relative to the number of steps that users have to perform to clean up Mirai [41, 55, 244] and Windows malware (e.g. running an antivirus). The difficulty of removing a QSnatch infection in users' home networks could be affected by these three pillars. If any of these attributes are not in place, this can translate into a longer time to remove the malware infection.

The COM-B model then stresses the importance of individuals' capabilities, motivations, and opportunities to perform a behavior. These aspects are critical for moving from malware detection to targeted intervention, and ultimately to user's actively adopting and proactively using malware-prevention solutions. Framed this way, the partner ISP was deploying an intervention, to notify users of the QSnatch infection and prompt a new behavior to occur. The COM-B model can help us to understand whether the COM attributes are being supported, and if any one pillar is not sufficiently supported, toward influencing ISP customers to perform a particular behavior. This behavior may or may not lead to the cleaning of infected devices, so we can recommend how the current intervention or future interventions can be improved. To add value to the partner ISP, COM-B is suitable for analyzing customer behavior after they receive the notification, to identify where targeted improvements may be made.

In reference to the COM-B model, we asked our participants a range of questions, addressing various aspects key to a successful behavior; the opportunity presented by the intervention from the ISP, in this case, receiving the notification (including whether it was noticed, and trusted, as in C.2); participants' capabilities to parse and act on the content of the notification (such as existing experience with IT systems and if users asked for help), and; if users had any limitations or reservations about performing the steps (such as perceiving a lack of support or tools to complete the steps in the notification, or beliefs about their own capacity or urgency to take personal action).

CODING AND QUALITATIVE DATA ANALYSIS

Once interviews were completed, they were transcribed and analyzed. Two of the researchers coded the transcripts with ATLAS.ti software using codebook-style Thematic Analysis (TA) [44]. Codes were created to label recurring topics, guided by discussion between the two coders to refine the themes. Inter-Rater Reliability (IRR) does not impact the usefulness of emerging themes with this approach, as noted by Braun & Clarke [44] and others [183]. However, themes were discussed at intervals with the wider co-author team to determine the central themes, where this approach can ensure the reliability of findings [183]. Agreement was reached on seven categories that pointed to core themes in subsection 4.4.2. The last theme on *Suggestions* was related to customer feedback, mostly as recommendations for improvements to the service. *Suggestions*, are then included in the Results section (section 4.4) where they relate to other core themes and not as a stand-alone subsection, more specifically in Table 4.4.2, and they were also shared with the partner ISP after concluding the research, to inform considerations for improvement to the support that the ISP gives to its customers (See subsection 4.3.4).

Table 4.1 shows an overview of the core themes, along with examples of codes within each core theme, and the percentage of respondents that discussed those themes as an indicator of the prevalence of each theme across the participant cohort.

Table 4.1: Summary of qualitative coding scheme

Themes	Code examples	Respondents n=57
Receiving and understanding the notification	Receiving notification, understanding notification message	57 (100%)
Cleanup effort	Cleanup time, time to execute steps	49 (86%)
Technical (security) ability	IT profession, IT experience	56 (98%)
Beliefs about risk	Consequences of not executing steps	54 (95%)
Responsibility	Personal responsibility, ownership	53 (93%)
Communication channel	Trust, distrust	41 (72%)
Suggestions	Suggestions to the ISP, comments to the ISP, congratulations	34 (60%)

PILOT INTERVIEWS

It was important to arrive at a robust study protocol, not only for engaging with real-world users outside of a controlled laboratory setting, but also with participants who were customers of our partner ISP. To test the interview protocol, 45 customers were contacted, 14 customers did not answer the call, 14 opted out, 3 numbers were out of service, and 14 customers decided to participate in the research. From the 14 customers who participated, 7 (50%) customers were showing up as remediated at the moment we talked to them and 7 (50%) were showing up as still infected. The main change after the pilot was to ask users if the device was used for private or business purposes or both. We uncovered that some customers

use their devices for these different purposes. The pilot interviews led us to decide to have more precise questions and less open questions. This was based on the willingness of ISP customers to participate in the pilot since we learned that customers would not spend on average more than fifteen minutes engaging with the data collection, this without including the time that the researcher carrying out the interview took to introduce himself, describe the research, and gain consent from the participant. These 14 pilots interviews are not included in the dataset of 57 interviews that forms the basis of the interview study.

INTERVIEWED PARTICIPANTS

After completing the pilot, we conducted 57 interviews. The age of these customers ranged from 22 to 63 years old. Four (7%) participants self-report their gender as female, and 53 (93%) as male. Most participants, 46 out of 57, used the QNAP device for private purposes, 5 used the device for business purposes and 6 used the device for both business and private purposes. No incentive was provided to participate in the research.

4.3.3. CLEANUP TIME AFTER NOTIFICATION

While we derive the infection time (or infection duration) from the Shadowserver data (subsection 4.3.1) recorded in the incident ticketing system of the ISP, we also want to know how long it took users to clean up the infection after they were notified. In this research, we defined as “clean” a user device which stops showing up as infected after being notified. On the other hand, if the observation continues showing up in the feed, we considered the user as “not clean”.

Retrieving the time stamps of the notification(s) was a labor-intensive manual process. Since the abuse department was transitioning to a new system, it was required to manually check the IDs for a period of a year and be careful about not missing notification. Thus, we were only able to do this for the interviewed users, except for six customers, where the abovementioned system transition meant we could not retrieve this data.

In the end, we collected the notification time stamps for 51 users. For this group, we could determine when the cleanup happened in relation to the notifications received by the customer. Unfortunately, we could not compare these findings with the Mirai and Windows infection groups.

4.3.4. ETHICAL CONSIDERATIONS

The human research ethics committee of our institution approved the interview protocol of this study (Reference number: 1490). Consent for anonymously taking part in this research, as well as for recording the calls, was obtained from the participants. They were also

reminded that they could stop the study at any time.

Following the Menlo Report [88], we were guided by the ethical principles of respecting people, respecting the law, justice, and beneficence. Regarding respecting people and law, we followed all the guidelines and privacy policies of our partner Internet Service Provider, and personal data never left the Internet Service Provider's premises. One author was embedded in the ISP, and in consultation with the ISP's privacy team and within terms of service linked user IPs and interviewees then produced an anonymized dataset used in our data analysis. Unfortunately, even though we used an anonymized dataset, our partner ISP did not agree to make the data publicly available. Further, as pilot participants stated that they had limited time to participate in research, the protocol for the main study was adapted to respect this.

Regarding justice, the study did not benefit specific groups over others. All infected customers were contacted for the study and had equal opportunity to share their experiences and provide feedback.

Regarding beneficence, we did not interfere with the ISP's beneficence and all subscribers affected by QSnatch malware were notified of the infection, so they were able to protect themselves and others from this threat. The goal of the interviews was to learn how users experienced the remediation process to improve the support that the ISP can give for its customers. Also, our research aims to understand how users deal with persistent IoT malware in order to benefit society at large.

4.4. RESULTS

4.4.1. SURVIVAL ANALYSIS

In this section, we answer the question of whether persistent IoT malware, namely QSnatch, is more difficult to remediate compared to persistent Windows malware and non-persistent IoT malware, namely Mirai. Higher difficulty would result in longer infection times.

Table 4.2 shows the cleanup success and the infection times (mean, standard deviation and the distribution) for each of the three categories of malware families, namely Windows malware (n=425), Mirai (n=107), and QSnatch (n=228).

The mean infection time of Windows malware is 36 days with a standard deviation of 76 days, the median infection time is 0 days and the maximum infection time is 359 days. For Mirai, the mean infection time is 19 days, with a standard deviation of 76 days, the median infection time is 1 day, and the maximum infection time is 182 days.

In contrast, for QSnatch the mean infection time was much longer than the other two malware categories: 108 days, with a standard deviation of 110 days. The median infection time is 76 days, and the maximum infection time is 365 days.

Table 4.2: Summary statistics per group of infection type, with remediation outcomes.

Group	Sample Size	% clean	Infection time (days)						
			Mean	Standard deviation	Min	25%	50%	75%	Max
Windows malware	425	97%	36	76	0	0	0	26	359
Mirai	107	100%	19	36	0	0	1	24	182
QSnatch	228	91%	108	110	0	3	76	181	365

Table 4.3: Summary statistics for interviewed and non-interviewed groups exhibiting QSnatch device infections.

Group	Sample Size	% clean	Infection time (days)						
			Mean	Standard deviation	Min	25%	50%	75%	Max
QSnatch – Not interviewed	173	89%	112	116	0	3	76	181	365
QSnatch – Interviewed	55*	100%	94	86	0	2	76	157	273

* Note that the interviewed group is $n=57$. We could not retrieve the infection end dates for two users, due to the system transitioning at the abuse department (See subsection 4.3.1), thus in this table $n=55$ for the interviewed group.

The mean infection time of QSnatch infections is three times higher than the mean time for Windows infections and five times higher than the mean time for Mirai infections.

For a more comprehensive analysis of the data, we computed the survival probability for each malware category using Kaplan-Meier estimates. Figure 4.2 shows that after 180 days, around 30% of the QSnatch infections are still alive, while only 10% of the Windows infections remain, and none of the Mirai infections. Figure 4.2 inset figure also shows that within 7 days after the infection QSnatch remain stable at almost 80%, while Mirai and Windows malware already drop to almost 50% or lower.

Consistent with [54], we have also observed a high cleanup rate at the beginning of the infection time for Windows malware and Mirai, even though in our study, most participants were notified via email rather than put in a quarantine network. We do not observe this same pattern for QSnatch.

The log-rank test reports whether there is a significant difference between the QSnatch and the two other groups. We find that the differences with both groups are highly significant: Mirai versus QSnatch (log-rank test: $X^2=96.22$ with $p=0.00$) and Windows malware versus QSnatch (log-rank test: $X^2=80.27$ with $p=0.00$).

To check whether our interview study suffered from selection bias, where the people who were willing to participate might also be more committed to conducting remediation, we




Figure 4.2: Survival probability QSnatch vs Window malware vs Mirai

analyzed the infection time data for both groups, interviewees as well as non-interviewees.

From the total users infected with QSnatch in the period of observation ($n=228$), we interviewed 57 users. We need to remind the reader that we could not retrieve data of 2 participants for the survival curve, thus the number of observations in the graph is 55 (See subsection 4.3.1).

Table 4.3 shows the summary statistics of the QSnatch not-interviewed users versus the interviewed users. The mean infection time of not-interviewed users is 112 days with a standard deviation of 116 days. For the interviewed group, the mean infection time is a bit lower, 94 days, with a standard deviation of 86 days.

Figure 4.3 shows the survival probability of both groups. They are very similar. Only at the tail end of the plot do we see that 10% of the non-interviewed group remains infected at the end of the period, while all of the interviewees have remediated. We did a log-rank test to check if there were differences between the groups. The log-rank test reports no statistical differences between the groups at a 5% significance level ($X^2= 3.09$ with p 0.08).




Figure 4.3: Survival probability participants vs not interviewed users.

4.4.2. INTERVIEWS

Given that participants have to comply with many different steps, as described in section 4.3, different questions were asked to understand how they handled the remediation process. The interviews were transcribed and coded, different themes emerged that will be described in this section. We focus on the most prominent themes which emerged from the interview analysis.

RECEIVING AND UNDERSTANDING THE NOTIFICATION

We asked participants if they recall receiving the notification and recall performing the recommended steps. In total, 53 (92%) recall receiving the notification, and only 4 (8%) participants either said they did not receive it or were unsure. Table 4.4 shows a summary of the participant's answers. The majority of participants 44 (77%) recalled doing the steps. Of the remaining 23%, most customers reported charting their own course to solve the infection. Five (9%) participants reported updating the device, while three (5%) reported following the manufacturer's steps, and four (7%) turned off the device. One (2%) respondent reported calling the QNAP helpdesk to solve the infection. The four participants who said that they did not receive the notification or were not sure of receiving it, were among the participants

that mentioned doing some of these different steps.

Table 4.4: Recall of responses to ISP notifications for interviewed participants (N = 57).

Recall of responses	Interviewee response	No. Interviewees
Recall receiving notification	Yes	53 (92%)
	No	2 (4%)
	Not sure	2 (4%)
Recall performing steps	Yes	44 (77%)
	Update	5 (9%)
	Turn off device	4 (7%)
	QNAP steps	3 (5%)
	Call QNAP helpdesk to cleanup	1 (2%)

CLEANUP EFFORT

We asked the participants to estimate the time they had spent on their remediation actions. There was high variability in the reported times and four categories emerged.

Table 4.5 shows the self-reported time participants invested following the steps. The largest group of participants, 25 (44%), gave answers in the range of more than 15 minutes up to 1 hour. Ten (18%) participants reported taking up to 15 minutes. Eleven (19%) reported answers that ranged from more than one hour up to twelve hours. For example, P45 reported *“The steps you indicated were completed quickly. That would have taken half an hour. But what actually should happen was that it took half a day of work to solve it completely.”* Finally, three (5 %) participants reported taking more than 12 hours up to 24 hours. In cases where participants reported almost a day of work, they did not refer to the actions themselves consuming so much time, but that their overall remediation process took that long. The NAS took time to execute various instructions as well. For instance, P47 stated *“I think (I spent) an hour per NAS myself, but the device can easily be working for an entire day”*. Eight (14%) participants did not answer the question. All in all, participants reported being able to execute the actions swiftly or at least within a day.

Next, we looked at the cleanup time: the time between the first notification and the end of the infection. We have the time stamps of all notifications for 51 participants. (As explained in subsection 4.3.3, we could not retrieve this information for six customers.)

The largest set of participants, 24 of 51 (47%), acted after the first notification. All of them cleaned up in one or two days after the notification. Note that some of these participants received their first and only notification very late, because of the random delays caused by the abuse system transition. If they were not immediately notified upon the first observation in the Shadowserver data, some time would pass before they are observed again in the Shadowserver data. In some cases, even this second or third observation did not trigger a

Table 4.5: Self-reported time invested in remediation actions

Time	No. Interviewees
Up to 15 min	10 (18%)
More than 15 min, up to 1 hr	25 (44%)
More than 1 hr, up to 12 hrs	11 (19%)
More than 12 hrs, up to 24 hrs	3 (5%)
No answer	8 (14%)

notification. This meant that their infection time could be very long. Because of the random delays, nine of these customers have a total infection time between 31 days and 241 days. Yet, once the customer is notified, the remediation takes place within two days at most.

These random delays unintentionally tested the effect of the notification. Before the notification, those users were infected with QSnatch for one or more months. It underlines the necessity of the ISP notification. Apparently, there is no alternative path towards remediation.

Next, there is a group of customers who did not act shortly after the first notification. The number of notifications that a customer received is clearly connected with the overall duration of the infection. 21 of 51 participants (41%) received between 2 and 9 notifications. For these customers, the duration of infection ranged between 8 days to 252 days. Finally, there were 6 of the 51 (11%) participants who were notified between 11 and 18 times. The duration of the infection was between 128 and 273 days.

From the interviews, we could identify reasons why participants did not act immediately on the notifications. We found evidence of users planning the remediation tasks consistent with [279] that might have delayed the action. P15 said he wanted to wait until he “*could take my time on a Sunday to try to solve this*”. Another interviewee, P8, referred to outsourcing the tasks. He hired an IT provider to do the steps and that process took a while. This participant received 11 notifications before finally showing up as cleaned. P19 said he received the notifications while being out of the country and without remote access to the device, so he had to wait until he got back. This participant received 11 notifications. Other participants said they did not see the email immediately, because it arrived in a mailbox that they do not frequently check. P52: “*The emails arrived early with me, but in a mailbox that I barely ever read. That is why I checked this only very late and took care of the situation.*”. This participant received 13 notifications. P20: “*It took a while before I saw [the notification]. Once I saw it, I took action*”. This participant received 18 notifications from the ISP.

P50 was one of the customers with the most notifications, 18 in total. When he was asked whether he found the steps useful. He said: “*Yes, although it is difficult to know whether it is useful. Initially, I ignored the email twice or so, because I wondered whether this was*

officially from the ISP. You get so many emails these days that you aren't sure. But after receiving it repeatedly, 2 or 3 times, I thought: OK, this is serious, let's take action."

The more participants overlooked the notifications, the longer their infection time. That said, the recurring notifications did at some point spur them into action. Only one customer cleaned up long after the last notification.

Our data shows that many participants acted on or close to the first notification they received. It is also important to note that email was an effective channel through which to reach many users, similar to [57]. Our findings demonstrate that email can be a cheap and scalable alternative to prompt participants to remediate, compared to walled gardens, letters or other notification mechanisms [178]. However, for some users an alternative notification mechanism, rather than a repeat notification, might be needed.

Finally, the random delays in the first notification also demonstrated how important the notification process is for persistent malware.

TECHNICAL ABILITY

The remediation process entails several relatively complicated steps, compared to the remediation advice for Windows and Mirai infections [41, 55, 244]. Yet, most participants report needing only a short time to conduct the steps. In line with this, we also encountered very little evidence that participants felt the steps were difficult to execute. Only one person mentioned any doubt as to how to perform to remediate the problem. Most participants described it as a straightforward task. This sounds a bit paradoxical: the task is relatively difficult, the infection took long, users were not aware of the infection until the ISP notified them, yet very few users expressed experiencing any difficulty.

This paradox points to their skill level related to capability in the COM pillars. We asked interviewees about their IT experience. Table 4.6 summarizes the answers across three main categories. A stunning 56% of the interviewees said they were IT professionals. P43 said: *"I've worked in IT for 20 years"*. Several participants said they worked as system administrators. For example, P44 mentioned: *"I'm kind of a system administrator at work. I'm fairly well versed in it"*. Others work in software development and programming—for instance, P56: *"I have my network in my home, you know, I can use it, this is also my job, my profession is programming"*. Some reported working in network security and automation.

A second group, 16% of the participants, claimed some experience managing IT, though generally out of interest rather than in a professional capacity.

To illustrate, P52 said that *"I happen to work at [ISP NAME] myself"*. Others mentioned working on their own networks at home as hobbies, managing their own servers, and similar activities. P28 reported: *"So I've always had a server running in my own network for 15 years. A hobby that got out of hand."* Only 15 out of 57 interviewees (26%) said they did not

Table 4.6: Self-reported IT experience

Type of experience	No. Interviewees
IT professional	32 (56%)
No experience with managing IT	15 (26%)
Some experience with managing IT	9 (16%)
No answer	1 (2%)

have any real experience with IT. Four (27%) of these customers reported looking for help from IT professionals, an acquaintance with IT experience, or a friend. It is worth noting that from the users who did report some IT experience, one asked for help as well, but the help he referred to was contacting the QNAP helpdesk.

4

Our findings suggest a self-selection process or early adopters at work [84]. NAS devices attract a user population that is significantly more skilled than the average user population. This would explain why the participants handling the remediation had some tolerance to execute the complex process. In fact, for an IT professional, the frame of reference is different. They are more likely to compare the QSnatch remediation actions to IT admin tasks, rather than to the consumer tasks of running an AV tool on a Windows machine or changing the password on a Mirai-infected IoT device. In that light, the QSnatch remediation process is not particularly difficult. Clearly, this finding is unlikely to hold for other IoT devices that are more widely distributed among consumers.

Interestingly, only a few participants questioned how the ISP knew about an infection that they did not know about, or about how they got infected given that they have self-reported IT experience. Meaning that they did not question their setup or how they got infected. P14 stated: *"I'm curious how that [infection] came about. I have to say that I thought it was strange because I suspected that I had nothing open. I had all those services turned off. I only use it as a local NAS in my local network. So all ports to the outside were turned off. And that makes it very strange that that is possible."* In total only five participants were doubting how the infection happened or how the ISP knew about their infection.

BELIEFS ABOUT RISK

We asked participants what they believed would happen if someone were not to follow the recommended steps. Participants expressed certain beliefs which contribute to their decisions about whether to act upon the notification. We found a variety of beliefs about viruses, comparable to those identified in other work examining home participants' mental models of security [288].

Most participants state that if the steps are not followed, malicious activity may be directed toward them. For instance, the malware stays, data is lost or held for ransom, or the device becomes accessible to attackers, among other beliefs. For example, 18 (31%), stated

that unless action is taken, the malware will stay on the device. Two of these participants added that this could bring consequences to their network safety, and two participants mentioned that this could affect others. One participant said that the malware could spread. 13 out of the 57 participants described data loss or theft as an anticipated consequence of not completing the steps; one of this same group also mentioned the possibility of a Distributed Denial of Service attack. For instance, P16 expressed: *“it may just be that they can access your photos, for example, and do something with them, ransom and so on”*. P29 stated that *“then it [the malware] releases files that may be private”*.

Six other participants (10%) described that a compromise of the data on the device could be possible or that the device is made openly accessible for attackers and exploits. Three participants believed that they would not have access to the device due to malware (which potentially contradicts their having use of it at the time). Three participants mentioned that they could lose their Internet connection. This can be associated with the fact that the ISP notification stated that if the respondent did not complete the steps, that there was then the possibility of temporarily placing their connection in quarantine. Three also stated varied beliefs like the ISP would get into problems, that they would get into problems for many years, or that not doing the steps was not an option. Three participants were unsure of what could happen.

Most of the expressed beliefs, similar to [114], were about how participants think the malware would affect them individually rather than thinking about how the infection could affect others.

Relating to the clean-up behaviors, we then see that our participants were completing the steps and motivated to do so. This demonstrates a close link between security beliefs and protective behaviors [289]; where Wash & Rader found that individuals with a strong belief that viruses caused problems then self-reported taking action to protect themselves, we have real-world evidence here of this being borne out for consumer IoT devices (independent of the accuracy of the belief).

RESPONSIBILITY

When asked, 53 out of 57 interviewees (93%) said they felt responsible for cleaning up the device. Most of them, 34, expressed that the device belongs to them, they manage it, it is in their own network, and they felt responsible for solving security issues. To illustrate, P1, stated: *“Yes (I am responsible), my children my wife use the NAS so it must all be safe and there are also so private things stored there also, like tax data”*. Five participants connect feeling responsible to being informed of the problem via the notification. To give an idea, P48 said: *“yes, (I am responsible) because I was asked and I manage that system at home. So then I am responsible for those steps”*. The rest of the participants expressed diverse

reasons why they felt responsible for doing the steps, either they indicated that they were the ones having the problem, that no one else would do it, or that they felt responsible because they wanted to get rid of the malware before it caused potential damage.

Beliefs about responsibility are important, as other research has found that individuals may otherwise defer or delegate responsibility to other people [94]. We did not see this with the majority of our participants, aside from the few who approached an outside IT specialist for help. Even this action can be seen as a form of taking responsibility.

Haney et al. [139] asked an open version of this same question to smart device owners, finding a mix of perceptions across personal, manufacturer, and government responsibility; the majority of their participants stated at least partial personal responsibility for the security of their devices. Interestingly, their participants focused on personal responsibility specifically around fixing lapses or precautionary measures around device security which may result in exposure to risks – this tallies with the setting of our study, where our participants are uniquely queried about real-world infections of their own smart devices, and expressed personal responsibility to resolve the issue.

COMMUNICATION CHANNEL

During the interview, participants were offered the opportunity to discuss or mention things that they considered important that were not asked by the interviewer. Some participants discussed the trust issues they had with the notification. In total 12 (21%) participants mentioned feeling some distrust towards the ISP notification. Where participants provided customer feedback, most of their suggestions about the service related to the communication channel.

P27 stated: *“Those messages from [ISP NAME] looked very much like it was all fake, so to speak. So I was a little unsure about that too.”*. Also, P33 mentioned *“The mail I received from [ISP NAME], I got it in the spam folder, so I almost deleted it. [...] I liked it, I think it’s a very nice initiative from [ISP NAME], but to say that it is very normal, no. So it would almost look like someone is trying to trick me about my device. So the ISP should communicate a little better about that.”* Several users recommended to make the communication more trustworthy.

The level of distrust is higher than reported in a previous study, where only two users distrusted the notification via email [55]. The higher level of distrust might reflect the technical ability of the NAS owners, compared to the broader user population in the earlier study. Another explanation could be that participants received the notification during the COVID-19 pandemic. They were working from home and might have been more careful with the emails they received. Consistent with [266], the trust issues around the email could have played a role in delaying the actions as well.

4.5. RELATED WORK

Before the past two decades, Windows malware has occupied the security community [72]. Also, some Windows malware, such as Conficker, remained in users' machines for many years [26]. With the proliferation of IoT devices, now attackers are shifting to IoT persistent malware [40, 47, 276, 287] since these devices have several advantages for attackers, like low computational capacity [162], thus they cannot count on protections such as antivirus which Windows systems do. The current state of the art has also learned about Mirai, a non-persistent IoT malware, that can be removed by rebooting the device and changing passwords [22, 41, 55, 97, 244], however, in this research we dealt with QSnatch, a malware, that needs convoluted steps from users to be remediated, and does not count on the same mechanisms for removal from Windows malware or Mirai.

Users were notified about a QSnatch malware infection in their home networks. Li et al. [173] studied notification content and mechanisms in terms of webmasters cleaning up compromised servers. They observed that contacting the webmasters directly increased the likelihood of cleanup by over 50%. In this study, we contacted the person who managed the network access storage device, and we observed that 45 (78%) of the participants did the recommended steps, and 13 (22%) charted their own course to solve the security problem.

Stock et al. [266] and Cetin et al. [57] sent notifications to vulnerable domains and described low remediation rates. They highlighted the limitations of email notifications and the breach between taking action and knowing about the problem. In our work all customers were notified via email only, and for some users the first email notification was enough to take action. However, some participants needed multiple notifications to act.

Li et al. [172] notified network operators about security issues in their networks revealing that different notifications have different outcomes, but in general notifications have a positive impact on remediation. Dumeric et al. [95] sent notifications for vulnerable Heartbleed servers and found a beneficial influence in patching. Cetin et al. [54] found high remediation rates for Windows-based malware cleanup. In this research, we observed total cleanup after participants being notified. This could be explained due to the capability that most users self-report.

Vasek et al. [282] studied how detailed notifications caused more remediation of compromised websites than short notifications. In this study, we found that users benefited from a tailor-made precise advice to execute the steps to solve QSnatch infection.

Different work on IoT malware notifications [41, 55, 244] highlight that once users are aware of an IoT malware infection, they are motivated, comply with the steps and cleanup. Our findings demonstrate that even with more convoluted steps users put time, effort and take responsibility to remediate the infection.

4.6. DISCUSSION

Our observations of network data illustrated that the mean infection time of persistent IoT malware is greater than that of Windows malware and memory-resident IoT malware; QSnatch infections may persist for several months, as also shown in our data. In terms of successes, we have found real-world evidence of our participants successfully mitigating persistent IoT malware. This demonstrates a close relationship with the intervention of the participants' ISP, where the QSnatch-infected devices of the customers we interviewed were remediated at a time close to having received a notification from the ISP. Issues arose in noticing one notification in a series of notifications, as the prompt to take action, and in subsequently planning to take action. In this section we discuss the wider implications of our quantitative and qualitative results.

4

4.6.1. SUCCESS AND TIMELINESS OF REMEDIATION

The participants we interviewed as part of this study all reported taking action to remediate; all were seen to no longer appear in the infection data shortly after receiving a notification. This implies that at least for participants such as ours, who believe they comprehend and can action advice when prompted, that this model of ISP notification is successful. Many participants were thankful for the notification.

No participant acted prior to receiving a notification, even if their infection was already going on for months. They did not report acting on unexpected device behavior before receiving the notification, as might happen with malware that is generally used to target others outside of the network. Given the proximity of a notification to remediation for participants, we posit that they may well have not taken action if they had not been notified. Natural remediation did not occur either (as has been noted can occur for non-persistent malware infections such as Mirai) [55].

We see from our results that, generally, those participants who took longer to remediate had received more notifications before eventually acting on one. It is less a question of whether we need to help people to successfully remediate, and starts becoming a question of whether we want them to remediate *sooner*. For researchers, this highlights the importance of combining self-reports with technical data, to understand where users are not noticing notifications compared to what they report [236]. This includes any contributing circumstances, such as seeing the notification when not being near the affected devices and being able to act on the advice (and forgetting it shortly after). In studying operating system warnings on personal computers, Krol et al. [165] found that over 80% of their participants were observed to ignore one or other warning, more than those with higher computing proficiency. Egelman et al. [98] found participants receiving a passive phishing warning mostly seemed to ignore

it, as compared to active warnings which require explicit interaction – email notifications follow a similar format.

4.6.2. LEARNING FROM THE IDEALIZED IoT USER

In a way, our study found an idealized version of ISP-managed remediation – the ISP has done what is within their power (acquire abuse data and send a notification) and our participants, for the most part, have received the notification, understood it, acted on it, and then their network is seen as being cleaned. There are, as mentioned above, some inconsistencies in that story, foremost that some participants required many notifications before acting (Opportunity).

This user population arguably consists of ‘early adopters’ of what is currently a niche device (network access storage devices), whose response to this emerging threat could inform what we can reasonably expect of users of varying expertise as this family of devices sees more widespread use. Foremost, this user group was relatively ‘cheap’ to help – they were told what to do (Opportunity), they did it (Motivation), and it worked. We cannot assume this would hold for other groups of smart device users, especially those with less technical experience (Capability).

The role of personal responsibility in keeping IoT devices secure has been highlighted in other work [139], and further, that suitably informed personal responsibility requires understanding of communicated risks, the opportunity to act, and to know how to act. We saw a few participants take independent action to verify the right steps to take, rather than follow the notification steps exactly (as in C.1). This suggests that less tech-savvy users may also need advice pitched at a suitable level of competence – a few of our successfully-remediated participants reported updating the device, or calling the manufacturer for support, which are both approaches which can be adapted to less technically-experienced device owners as they rely less on an assumption of personal technical ability.

4.6.3. SELF-EFFICACY AND DEVICE COMPROMISE

Our participants seemed confident of their capacity to clean up the devices. Despite having been informed that their devices were compromised, none mentioned being surprised or doubting the correctness of their device setup. However, five participants did enquire as to how the ISP knew about the infection and, in a manner, questioning whether there was an infection. It may be that our experienced participants do not associate remediation efficacy with device setup efficacy. Otherwise, the IT-related work that many were involved in may have desensitized them to device infections being an issue, especially if they perceive it as not directly affecting them personally, and hence some lack a sense of urgency (Table 4.4.2).

Our participants then bear resemblance to users who are ‘engaged’ by security [126], who when alerted to there being a problem will seek out a solution.

4.6.4. LIMITATIONS AND FUTURE WORK

A limitation of our methodology is that study was carried out in a single Internet Service Provider (ISP) that has an established process for notifying users. Thus more research might be necessary to compare how this process happens in different ISPs. Additionally, we focus on a single persistent IoT malware family as a case study, QSnatch. Other persistent malware families might require different steps, and they might be harder to remove [40, 47]. Thus, future research could consider comparing different persistent IoT malware families, to understand the applicability of our findings to other cases.

As with previous work [41, 244], the results of this research were based on users self-reported behavior during interviews. The interviews were performed after a certain period of time. Ideally, we would have contacted users close to the notification time. However, the ISP was already notifying customers as part of their abuse handling process as explained in the section 4.3. One of the authors was embedded in the ISP for a period of time; thus, we used historical data that allowed us to observe the QSnatch malware behavior over the period of a year. By using a larger historical timeframe, we could include a larger sample of affected users at the cost of more time between the remediation and the interview. If we wanted to time the interview close to the notification, then we would have to accept a much smaller user sample. In our results we found that 92% of the participants stated that they received the notification, and all participants recalled what they did with it; thus, there is no evidence that they forgot what they did. This is in line with earlier work. Studies of security experiences, such as software updates [281] and social diffusion of security-related behaviors [77, 78], gathered insights on user behaviors across potentially far-reaching timescales.

Finally, due to the manual intense process of retrieving the notification(s) dates, we could not compare the cleanup time of QSnatch with Mirai and QSnatch, thus future research could look into that. We only interviewed QSnatch infected users, thus we learned about their process of handling the infection. Previous work [41, 55, 244] has looked into the remediation process of Mirai, thus we focus on a group that has never been studied before, victims of persistent malware.

4.6.5. RECOMMENDATIONS

From our analysis, we provide the following Recommendations:

- **Adaptive notification channels.** An approach would be to find a manageable way to ‘ramp up’ successive notifications to users at scale. However, any additional

effort to encourage remediation across a sequence of notifications is borne by the ISP (who is already the stakeholder ‘taking charge’ of the problem for lack of direct engagement by manufacturers). In many cases, our data shows that participants acted in effect immediately upon seeing ‘a’ notification, albeit the last in a series of similar notifications. One approach might then be to consider other channels after the first notification (as seen in [58]).

Figure 4.2 also indicates that there may be diminishing returns for solely email notification (the QSnatch curve flattens out as time goes on). Our data also showed that many participants acted on or close to the first notification they received. Email was an appropriate channel through which to reach many, but some users may need an alternative notification rather than a repeat notification. Email as a notification channel works for some, but alternatives should be explored (within cost-effectiveness for an ISP), for instance, quarantining the connection, voice mail, direct phone call or letter to the customer.

- **Framing and planning within remediation notifications.** All but one of our participants acted on a notification that they received. There were issues for several participants in terms of deciding to act on a notification and then having to find an opportunity to enact the instructions. The notifications we studied here act as a reminder to imply that immediate action is needed, primarily due to evidence of an active malware infection.

A balance may be struck between this and the use of commitment devices in reducing postponement (as explored elsewhere for security update behaviors [127]), for example having a user set a reminder for themselves for the same evening or the next day. Framing is also important, where a few participants presumed the email notification was fake at first. This relates to messenger effects in effective communication of ideal security behaviors [48].

- **Tailor-made advice.** Advice was specific to QNAP, and specific to QSnatch infection, rather than requiring a diagnostic analysis to determine which steps to selectively apply, as per the recommendation of the manufacturer. It was thereby actionable, from the users who did not have IT experience ($n=15$), 9 (60%) reported following the advice (5 reported asking for help), and one user reported following QNAP steps.

Consistent with what [235] recommended, we found evidence that minimum and concise instructions work, when measured via the remediation of IoT malware infection. Most of the time, Internet Service Providers (ISP) are restricted by laws and regulations, such as the General Data Protection Regulation (GDPR), from collecting

data on the population of user devices in the local network. So in many cases, they cannot know in advance which is the infected device to provide tailor-made advice, thus we have to rely on generic advice for most cases. An intermediate point can be gaining consent from users to actually identify the infected device in their network to offer accurate help.

4.7. CONCLUSION

Internet Services Providers use different methods to communicate with infected subscribers, and according to best practices [176], email notifications is one of them. This paper shows that notifications play a crucial role in the cleanup process of persistent malware like QSnatch. In contrast to previous predominant malware families (e.g., Mirai or PC malware), an automatic scan or an AV tool or power cycle do not get rid of the QSnatch malware. As we observed in Figure 4.2 QSnatch takes a longer time to get clean.

Does QSnatch take longer to clean because it is hard? The remediation advice of the ISP consisted of a convoluted series of steps, however, most users reported having high technical competency. Participants did not find following the steps as a problem, there might be a self-selection process of users with some IT experience, so these users might be comparing the steps to IT tasks. Hence, we dealt with the idealized user, they are capable, they are motivated, but it clearly takes time to organize cleaning of the device(s), and the majority of users had to receive multiple notifications to prompt them to act. The necessity of an external prompt for them to act contributes the non-trivial nature of QSnatch remediation.

In this study, we found that there is a lack of feedback loop about infections and cleanup success. Users had to be notified in order to act. An external party, in this case, the ISP, had to tell users they are infected and provide tailor-made advice to execute the right steps. This is not always possible for the ISPs since they cannot know in advance which malware and which device has been infected.

Nevertheless, our study shows that all users remediated at some point, so damage is less with this self-selected user. It can happen that this will fall apart when average users use products affected by persistent malware, but it can also be that manufacturers such as QNAP are building already tools similar to Windows tools such as malware scanners and automatic updates from which average users will benefit.

In this study we have also found out, similar to [57], that email notifications could be effective. During circumstances such as a global pandemic where users depend on their Internet connection, this can be a good alternative. Similar to [41, 244] we found that when users are informed of a security issue, they are willing to act. They take responsibility although they do need some time and effort to execute the steps. In this, however, we

need to take into account that technical abilities are key, to comprehending the notification, understanding what needs to be done, and knowing how to do it in a sufficiently complete and error-free manner. Connecting this thread of interdependent activities was not difficult for our participants, but it may be for those who are less tech-savvy. This is especially important in the absence of direct indicators from smart devices as to their security status, as explained in the opening arguments of this paper.

Regarding future work, we found that participants who took a long time to remediate their devices had generally received the highest number of successive notifications from the ISP. More correlation of technical and qualitative data is required to understand the role of communications and communication channels, and users planning strategies, especially as persistent malware continues to affect consumer devices.

5

IoT MANUFACTURERS' ROLE IN DEVICE INFECTIONS

The influx of insecure IoT devices into the consumer market can only be stemmed if manufacturers adopt more secure practices. It is unlikely that this will happen without government involvement. Developing effective regulation takes years. In the meantime, governments have an urgent need to engage manufacturers directly to stop the damage from getting worse. The problem is that there are many thousands of companies that produce IoT devices. Where to start? In this paper, we focus on identifying the most urgent class: the manufacturers of IoT devices that get compromised in the wild. To identify the manufacturers of infected IoT, we conducted active scanning of Mirai-infected devices. Over a period of 2 months, we collected Web-UI images and banners to identify device types and manufacturers. We identified 31,950 infected IoT devices in 68 countries produced by 70 unique manufacturers. We found that 9 vendors share almost 50% of the infections. This pattern is remarkably consistent across countries, notwithstanding the enormous variety of devices across markets. In terms of supporting customers, 53% of the 70 identified manufacturers offer firmware or software downloads on their websites, 43% provide some password changing procedure, and 26% of the manufacturers offer some advice to protect devices from attacks. Our findings suggest that targeting a small number of manufacturers can have a major impact on overall IoT security and that governments can join forces in these efforts, as they are often confronted with the same manufacturers.

5.1. INTRODUCTION

Insecure Internet-of-Things (IoT) devices are still flooding the market, even though the damage that these devices can cause has been evident for years. In response, many governments have issued baseline security recommendations and guidelines for security by design for IoT [104, 115, 116, 150]. While useful, such guidelines do not address the underlying root cause: many manufacturers lack the incentives to adequately secure their devices. Similarly, engaging with certain other actors, for instance ISPs who are in a position to mitigate part of this problem on a short-term basis, still leaves much to be addressed [24]. A consensus is emerging that governmental interventions are required to overcome the incentive problem [250]. While governments are debating long-term solutions regulatory strategies like apportioning liability and setting minimum security standards, and some liability frameworks are being proposed [195], there is a short-term need to engage manufacturers to reduce the current influx of insecure devices. To illustrate: in 2016, the U.S. Federal Trade Commission (FTC) lodged a complaint against ASUS because the company “failed to take reasonable steps to secure the software on its routers”. Through a consent order, the FTC got ASUS to “establish and maintain a comprehensive security program subject to independent audits for the next 20 years” [118].

For governments, the process to engage manufacturers directly is resource intensive and can only be applied to a limited set. Where to start? The question of which manufacturers to engage is complicated by the enormous complexity of the IoT ecosystem. There are markets around many different product types, each with different populations of manufacturers. Kumar *et al.* [166] found a long tail of 14,000 companies, though just 100 of them were responsible for 90% of devices in their observations. There are geographical factors at play also. Product types and manufacturers will vary across different countries and continents. Last, but not least, governments lack reliable data on the security practices of these manufacturers.

As part of a collaboration with the Dutch government, this paper presents an empirical approach to identify the priority targets for governmental intervention: the manufacturers of IoT devices that get compromised in the wild. For those manufacturers, the evidence for the lack of adequate security of their devices is compelling, as is the fact that this lack is causing harm. To identify device types and manufacturers, we build on recent advances in large-scale device discovery and identification. As a basis for governmental action, though, these studies have certain drawbacks that we need to overcome. Some studies rely on privileged access to internal network data from home [166, 301] or ISP networks [218, 247]. For our purpose, this approach would create selection bias towards the manufacturers in the limited set of networks where such access could be obtained. Other studies use Internet-wide scans,

typically focused on developing scalable methods for identification [299]. For scanning, most studies use device fingerprints that were developed for a specific set of devices that the researchers had access to. In other words, these scans can only detect a sample of devices that the researchers knew about and could test beforehand. It is unknown how these samples relate to the population of devices in the wild. This means that all other IoT devices are simply out of scope. In our case, however, we need to identify manufacturers for a specific population of compromised devices in the wild, not for a set of devices that was predefined. The population in the wild will at best partially overlap with those discovered in the large fingerprint-based studies. Finally, most studies focus on IoT devices in general, not on compromised devices. The few exceptions have serious limitations; one is based on an observation period of a single day [201], one only identifies high-level device categories [23], and one relies on third parties such as Shodan [258] to identify manufacturers [130].

Our approach starts with two months of real-time observations of the IP addresses of compromised devices that were scanning a /16 darknet with the Mirai fingerprint. As most compromised IoT resides in consumer networks [55], we focused our analysis on devices in 355 ISP networks that together have the bulk of the market share in 68 countries. Each real-time observation was immediately followed up with an active scan of that IP address that collected banners and Web-UI pages for the device. We then manually labelled the unique device fingerprints in an attempt to identify as many manufacturers and devices as possible. We opted for manual labelling because our goal is to provide data that is as accurate, explainable and complete as possible, since it will provide the basis for regulatory interventions. The goal was not to improve on existing scalable identification techniques. Our approach was able to identify 31,950 compromised devices attributed to 70 manufacturers. We aim to answer these questions: (i) Which manufacturers are associated with compromised IoT across 68 countries? (ii) How variable is the set of manufacturers across different countries? (iii) What are these manufacturers doing to remediate the insecurity of their devices? In sum, we make the following contributions:

- We present the first systematic analysis of which manufacturers share attributed infections for infected IoT devices in 68 countries.
- We develop a transparent and reproducible approach to identify manufacturers of infected devices that can be applied across jurisdictions and that does not rely on privileged access to network data.
- Our results demonstrate a strong pattern of concentration: while we find 70 manufacturers in total, just 9 of them share around 50% of all infections. This pattern is quite consistent across multiple jurisdictions, thus supporting international regulatory

collaboration in engaging these manufacturers.

- Notwithstanding the variety across markets, geographical areas and legal frameworks, the set of manufacturers associated with infected devices is remarkably consistent across countries. The manufacturers related to around half of the share attributed to infections were present in at least 47 (69%) of the 68 countries.
- We analyze what, if any, firmware or software was provided to download by the manufacturer, and we found that out of the 70 manufactures 53% offer firmware or software to download on their websites. We checked if the manufacturers provide any password changing procedure, and 43% of them do. Finally, we checked whether or not there was some advice to protect the devices from attacks, and 26% of the manufacturers offer advice to protect the devices.

5

5.2. CONTEXT

IoT manufacturers continue to bring devices into the market at an incredible pace [203]. Many governments want to unleash the potential of this technology—e.g., the European Union (EU) highlighted IoT in its vision of the digital single market [106].

In light of the security issues associated with IoT, the engineering community keeps working on defining security standards. In 2019, the IETF published RFC8520 (Manufacturer Usage Description (MUD) [167]) aiming at providing a white list of their devices' traffic so third parties such as ISPs could identify anomalous traffic flows that do not match the MUD profile [240]. Governments have also acknowledged the need to intervene and define common guidelines to secure IoT devices.

On the side of governments, various countries are trying to change the behavior of firms in the IoT markets. The UK government released guidelines of what they consider a secure IoT product [150]. At the European level, the EU Agency for Cybersecurity (ENISA) has released good practices for secure IoT software development [104]. In the United States, the National Telecommunications and Information Administration (NTIA) created a bill to increase the transparency of the whole supply chain of IoT devices by encouraging the “Software Bill of Materials” (SBOM) [5]. In addition, the National Institute of Standards and Technology (NIST) created a inter-agency report (NISTIR 8259) to help manufacturers incorporate security into their IoT devices. NISTIR 8259 renders guidance on how manufacturers could provide post-sale security of IoT devices and on how to communicate security to customers [115]. NISTIR 8259A [116] asserts a baseline of security that an IoT device needs to provide through technical means. These government efforts could, in the long run, result in more secure IoT devices. Similar to the E.U. General Data Protection Regulation,

which is increasingly considered the default global standard for privacy [33], these IoT security policies might get manufacturers to follow them in all the countries where they have presence, rather than differentiate devices per jurisdiction.

In the Netherlands, there are discussions about an update-obligation law for 2021, which would make sellers of IoT devices responsible for supplying updates, rather than directly imposing this obligation on manufacturers [211]. Ahead of European and national regulation, the Dutch government—more precisely: the Ministry of Economic Affairs and Climate—wants to start conversations with manufacturers of poorly-secured devices to improve IoT security [187]. Our study is conducted in collaboration with the ministry and meant to provide the basis for the selection of manufacturers that the government will engage with.

5.3. ETHICAL CONSIDERATIONS

To answer our research questions, we deployed active scanning of IP addresses where we detected a device infected with Mirai malware. Since active scanning has ethical implications, especially when conducted in consumer broadband networks, we got the approval of the board of ethics of our university to start with this research (Application #993). Our Data Protection Impact Assessment (DPIA) and a Data Management Plan (DMP) were also reviewed and approved. we briefly discuss relevant ethical considerations organized around the principles laid out in the Menlo Report [88].

First, *Respect for Persons*. Since we cannot identify the owners of the devices located at the IP address that we scan, let alone being able to contact them, we cannot get their prior consent. On the IP address of the server conducting the active scans, we set up a web page with information about the project and an opt-out mechanism. We received four opt-out requests during the scanning period, and we removed these IP addresses from our dataset and from further scans.

Second, *Beneficence*. An unintended harm is that in a rare number of cases the screenshots captured from the scans would contain sensitive data, such as a customized NAS access login page that contained a personal picture. The data of all scans were stored on a secure server with access limited to the researcher team. The raw data was removed after the analysis was completed. The benefit of this research to the owners of the devices that we scanned back is that our findings regarding the manufacturers are part of a governmental project that aim to get manufacturers to better support these – and all other – users with insecure devices, a longer-term benefit that is underlined by the presence of a compromised device on the home network of the users involved in the scans.

Third, *Justice*. The selection of IP addresses to be scanned was driven completely by the observation of Mirai scanning traffic originating from these addresses towards the darknet.

Within this set, additional selection was made by focusing on IP addresses from broadband consumer networks. This process does not bias against specific user groups within the consumer population.

Fourth, *Respect for Law and Public Interest*. This study is co-funded by the central government and designed in partnership with them. It is part of the government 'Roadmap for Secure Hardware and Software' [187].

5.4. METHODOLOGY

To identify which manufacturers share attributed infections of the bulk of the compromised IoT devices in each country, we setup a (near) real-time data collection pipeline to gather information on infected IoT devices observed in the wild. This pipeline ran for a period of two months (July to September 2020). Subsequent steps were executed to process the pipeline data and arrive at a labeled data set of compromised devices and their manufactures. Figure 5.1 illustrates a high-level overview of our methodology. Steps 1 and 2 capture the real-time data collection pipeline. Steps 3 and 4 consist of the subsequent data processing and labeling components which were executed offline at a deferred time. Below, we provide more details on each step.

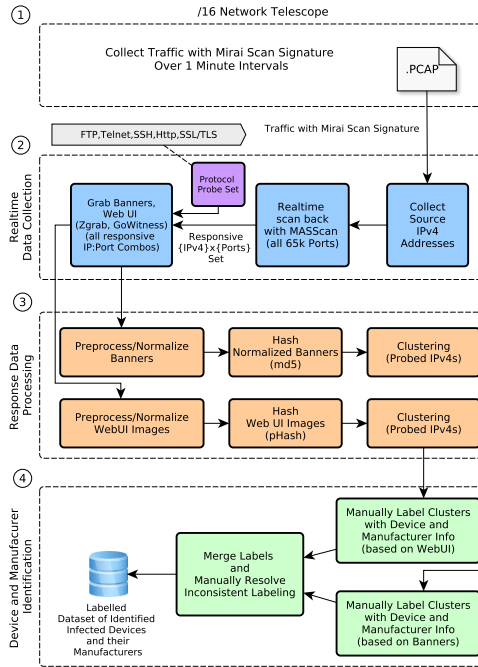


Figure 5.1: Data collection and processing pipeline

DATA COLLECTION

To collect data on infected IoT devices, we implemented a data pipeline tied to a /16 darknet through which we gather real-time observations on IPv4 addresses scanning the darknet with a Mirai malware fingerprint. We match all incoming darknet packets against the fingerprint developed in prior work on detecting Mirai [23] to filter and extract Mirai scan traffic from our darknet. All matching packets are buffered over 1 minute intervals and stored as PCAP network packet capture files which are queued for further processing in our pipeline.

Next – in step 2 of the pipeline – we extract source IP addresses from the queued PCAP files, and scan back all source IPs in (near) real-time to gather additional information on each entry. The data gathered here includes the set of responsive TCP ports at each IP, protocol banners for a set of pre-selected TCP services common to IoT devices (FTP, Telnet, SSH, HTTP(s), SSL/TLS), as well as screenshots of Web-UI content if publicly reachable through any of the exposed ports. This additional data helps us determine whether we are scanning back and potentially talking to a single device or multiple devices, and is simultaneously used to identify the IoT device(s) behind each IP and their manufacturers in later steps.

To gather this information, we first use `Masscan` [242] – a highly scalable TCP port scanner – to detect all open and responsive ports on the IP addresses in our data pipeline. We then feed its output to `zgrab` [2] to collect protocol banners for the previously mentioned set of common TCP services. We also feed the `Masscan` output to custom scripts to rapidly detect HTTP content on any responsive port whose output is then, in turn, used to collect screenshots of Web-UI content using `gowitenss` [1] a scalable Web-UI content collector implemented in the Go language. Note that we also probe for services on non-standard ports. Prior work has already demonstrated that the number of services running on non-standard ports are far more substantial than commonly assumed [149].

As a result of the large number of possible non-standard port and service combinations that need to be scanned, our pipeline has been tuned and highly optimized for collecting data expeditiously with all non-critical processing (and analysis) of the collected data deferred to subsequent steps. To further complicate matters, it is also crucial to maintain a near real-time scan back throughput in our pipeline due to potential IP churn. As IP addresses churn over time, the correspondences between IPv4 addresses and the IoT devices associated with each IP address will also change. An IP that previously corresponded to a network camera, now points to a home router for instance.

To maintain the necessary high scan back throughput we have implemented two main optimizations within the pipeline: First, we designed our pipeline to avoid scanning back an IP address that has already been scanned within the past 24 hours. We are assuming here that most IP addresses will churn at a rate slower than 24 hours. A secondary reason

for this optimization is ethical as we want to avoid directing unnecessary scan traffic to IP addresses and devices that have already been scanned recently. Our second optimization is due to observing a handful of IP addresses in our data pipeline that had all 65k ports exposed. We suspect these IP addresses to have pointed at improperly configured honeypots rather than actual infected IoT devices. We also observed a handful of IP addresses with an unusually high number of exposed ports. As a result, we optimized our pipeline to only grab banners and screenshots from the standard ports "21", "23", "80", "8080", "8081", "443" in combination with "FTP", "telnet", "HTTP", and/or "SSH" services when running into any IP address with more than 1,000 exposed ports after having scanned them via Masscan.

Note that a limitation of our approach – as well as all prior studies that employ comparable techniques to detect infected IoT devices from outside networks – is that an IP address does not have a one-to-one correspondence with a uniquely identifiable IoT device. With respect to the cardinality of the correspondence two corner case scenarios are possible in our case: (i) that multiple Mirai infected devices appear as having a single IP address due to Network Address Translation (for instance when multiple infected devices are sitting behind a router) (ii) that the infected IoT device is itself a router hosting an arbitrary number of other clean or infected IoT devices behind its NAT. With respect to these cases, we have adopted the following procedure: if the only device that we see accessible through our collected scan back data was a router and that router is known to be vulnerable to Mirai infection vectors, then we consider the router as the infected device. On the other hand if multiple devices have been detected, as long as they have known vulnerabilities to Mirai, all are considered infected.

In total, we scanned back 4,873,430 IP addresses using our pipeline. From this set, we selected the subset located in broadband ISPs for analysis. For these networks, we can have the highest confidence that the devices that scanned the darknet are actual consumer IoT devices, rather than scanners or other systems. Prior work also found that the overwhelming majority of compromised devices are located in broadband ISPs [55]. We used a reliable dataset of the Autonomous Systems (ASes) operated by broadband providers in 68 countries, developed in prior work [3, 24, 29, 177, 210], to filter and select the aforementioned subset.

Selecting for IP addresses in these networks, we had a set of 61,154 unique IP addresses. After removing results that consisted only of errors, such as 404, 401, we had a dataset of banners and Web-UI screenshots for 59,657 unique IP addresses.

For some devices, we could collect only banner data, but no screenshots of Web-UIs. For others, it was the reverse.

PROCESSING AND LABELLING

With the data obtained from the pipeline, we later performed two processes in parallel: (i) classification of banners, and (ii) classification of Web-UI images as depicted in steps 3 and 4 in Figure 5.1.

We first normalized the collected Web-UI images and banners. Similar to [271], we created heuristics with regular expressions to replace values such as date and time information, content-length field of HTTP response until the banners of the same devices were aggregated. We then hashed the images and banners. In the case of the images, we used perceptual hashing [153] to allow for minor variations in the Web-UIs. The banners were hashed via the MD5 algorithm. We then clustered the results based on the hash values.

Finally, we manually labeled the resulting unique 3,547 Web-UI image hashes and 566 banner hashes in our dataset and applied the labels to our clusters of data. These labels included manufacturer and device type. We identified manufacturers and device types based on the logos of the Web-UI and the text present in the banner. Hence, this labeling approach does not include original equipment manufacturer (OEM) because we observe the name of the brand that is marketing the devices. Regarding the device type, this approach sometimes allows to determine the category of the device (i.e. IP camera) and in some cases, it allows to get the specific model of the device. Finally, we resolved inconsistencies between the banners and Web-UI labels and we obtained our final labeled data set.

For about half of the devices, the collected responses did not contain any information from which we were able to identify the manufacturer. These results were labelled as ‘unknown’. From the 59,657 IP addresses, we managed to apply an informative label for 31,231 (52.3%) of them, corresponding to 31,950 devices. In total, we labelled devices for 70 unique manufacturers. As shown in Figure 5.2, we could identify 49 manufacturers via the banner data and 21 via the Web-UI images. There was an overlap of only 13 manufacturers. This underlines that any identification method would need to combine various types of data.




Figure 5.2: Number of unique manufacturers identified per data type

DATA COLLECTION ON FIRMWARE AVAILABLE AND MANUFACTURER SECURITY ADVICE

After we compiled a set of manufacturers and devices during steps 1-4, we also investigated what remediation options or security advice was being offered by the manufacturers. More specifically, we collected data for three categorical variables: (i) whether or not there was a software or firmware to download for the device model or device category; (ii) whether or not there was information provided on how to change the password for the device model or device category; and (iii) whether or not there was any security related information to protect the device model or device category from attacks. We followed an approach similar to [39], where researchers analyzed how security features and advice were presented to users in the manuals and support pages for 220 IoT devices. First, we identified the manufacturer's website. Since our approach sometimes allows to determine the category of the device (e.g. IP camera) and in some cases, it allow us to determine the model of the device (e.g. RT-AC5300), to accomplish this, we used Google's search engine with the following terms: "Device category" AND "Manufacturer" (e.g. IP camera Avtech) or "Device model" AND "Manufacturer" (e.g. RT-AC5300 ASUS). From the Google results we identified the manufacturer's website, which typically contains the manufacturer name in the domain name. Next, within the website, we manually inspect for "Device category" AND "manual" or "guide" or "quick start" or "Device model" AND "manual" or "guide" or "quick start" (depending on whether we had obtained the device category or the model) to check if the device model or category of the device had a user manual available. In cases where the search in the manufacturer's website was not fruitful, we used Google's search engine with the following terms to find the manuals: "Device category" AND "manual" or "guide" or "quick start" AND "Manufacturer" (e.g. IP camera manual Avtech) or "Device model" AND "manual" or "guide" or "quick start" AND "Manufacturer" (e.g. RT-AC5300 manual ASUS), depending on whether we had obtained the device model or the device category. In cases, where we had only the "Device category" (e.g. "IP camera"), we picked one random device of the category.

Next, we manually inspected for the "Device category" or "Device model", and we checked if there was a firmware (FW) or software (SW) to download available in the website. In the manual, we checked if there was any information on how to enable automatic "firmware upgrade" or "firmware update". The outcome was coded as yes or no depending on whether or not we found any FW/SW to download available either in the website or if we found any way to do automatic firmware upgrade or update in the manual. Next, within the documentation related to the "Device category" or "Device model" on the website or in the manual, we searched for the word 'password' to find whether the material contained any

password change procedure for the user of the device. The outcome was coded as yes or no depending on if a password procedure was found or not. Finally, we searched for ‘security’ as a keyword to inspect whether there was any information related to how to protect the device from attacks or make it more secure. The outcome once again was coded as yes or no.

To code our data, two researchers independently visited each manufacturer website and the manuals. Once they coded the three outcomes, they resolved inconsistencies by double checking the website and the manuals together. Figure 5.3 summarizes the method to check manuals and websites.

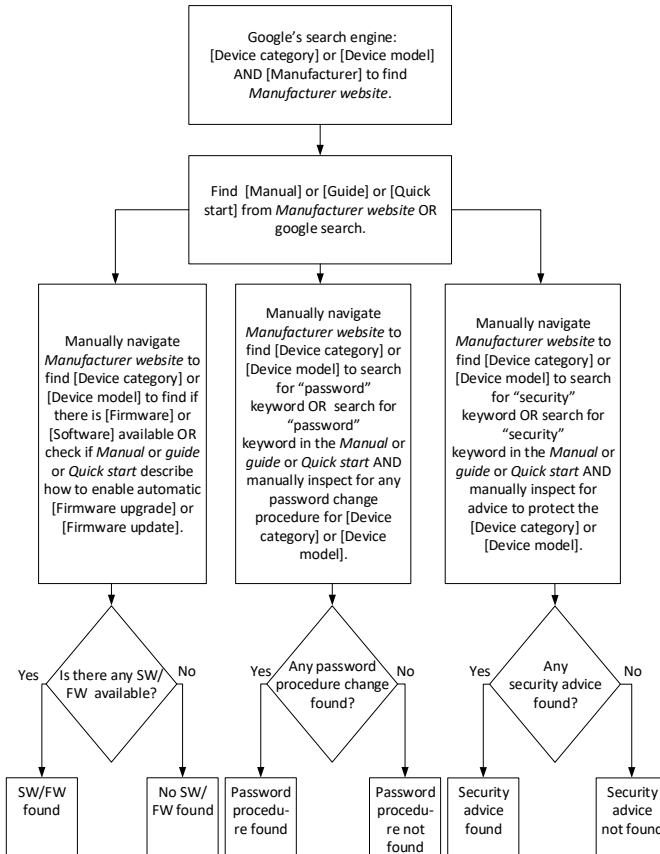


Figure 5.3: Method to check manuals and websites

Table 5.1: Average number of infected IoT devices seen per day for the top 20 countries (July–September 2020)

Country	Infected Devices (unique daily)	Subscribers	Infected devices (per 100k subs.)
Vietnam	10856	11959829	91
Taiwan	6627	4417500	150
China	6363	352767000	2
Rusia	2573	24125823	11
Brazil	2388	23529853	10
Indonesia	2240	7100350	32
Thailand	2183	8463797	26
United States	2136	94085580	2
Korea	1592	19073673	8
Turkey	1442	12159767	12
Mexico	1247	17432549	7
Italy	1227	16201874	7
Malaysia	1209	2492325	49
Iran	1092	10230000	11
Greece	1035	3912680	26
Egypt	1006	5133000	20
Romania	962	4513930	21
Germany	900	30868800	3
France	855	27292191	3
India	598	16410909	4
Others	13292	226277100	4

5.5. FINDINGS

Our final dataset contains data on infected devices located in 68 countries and attributed to 70 unique manufacturers—or labelled as ‘unknown’, where we could not identify the manufacturer. The number of devices seen in each country is highly variable. Table 5.1 depicts data for the top 20 countries with most infected devices.

Similar to [201] and [23], we find countries such as China, Vietnam, Brazil and the United States leading the number of infections. This suggests that number of infections is correlated to the number of broadband connections. This makes intuitive sense: more broadband subscribers means more devices connected to their networks, thus a higher risk of infections. To get a sense of the relative size of the number of devices in relation to the number of broadband subscribers, we have also included those statistics. We used Teleography data [6] of the first quarter of 2018 to calculate the total number of subscribers of the Internet Service Providers in each country. The last column contains the number of infected devices per 100,000 subscribers. There we see that countries with a large consumer broadband base have lower infection rates compared to many smaller countries.

5.5.1. MANUFACTURERS

Which manufacturers are responsible for the largest share of IoT infections in each country? Figure 5.4 shows that around 42% of the infections can be attributed to just nine manufacturers. Around 9% is attributed to all 61 other manufacturers combined ('Others (61 Manuf)'). We decided to group the 61 manufacturers because they were a long tail with a low share of the infections. The remainder consists of devices we could not attribute ('unknown').




Figure 5.4: Top 10 Manufacturers over the period of observation

There is a significant percentage of unknown manufacturers in our data. Could this potentially change the pattern of concentration that we found? In other words, could other major manufacturers be present in that set of unknown devices? To explore this issue, we looked at the frequency of the hashes. If hashes are seen only rarely, then it is very unlikely that these devices—and thus their manufacturers—make up a significant share of the population. (They could, of course, belong to one of manufacturers that we already identified. This would not change the overall picture, though, because of the small numbers involved). A high frequency for specific hashes, on the other hand, could point to the presence of a large share for a manufacturer.

We found that 57 unique hashes corresponding to banners and mainly for that part of the data we could not obtain Web-UI that provides information to allow us labeling the manufacturer either. We had 4 hashes corresponding to Web-UI. We plotted the cumulative probability of the number of hashes and less than 5 hashes have around 85% probability of showing more often (see Figure 5.5).

Some banners provided information about the device, mainly "DVR" and "NAS", but no manufacturer information. There were 17 (30%) hashes that correspond to DVRs and 6 (10%) to NASes. For instance, we got responses like 220 NAS FTP server ready. We are confident this is an IoT device, but it is not possible to determine with this information the manufacturer name. The rest was a long tail that included FTP servers and Bftpd




Figure 5.5: Frequency of hashes of non-identified devices

5

servers probably used by NASes, set-top boxes, and routers, however, it was not possible to determine the manufacturer either. We tried checking with different sources to determine if specific banner texts are unique to some manufacturers, but we did not succeed. Although this is a limitation of this method, which we discuss more in section 5.8, the frequency of the hashes gives us some confidence that the remainder of manufacturers that we could not identify will not change the overall picture.

The devices of nine manufacturers were responsible for a large share of the global set of infections. How dominant is this pattern at the level of countries? In other words, does each government have to engage with a different set of manufacturers or does the pattern of concentration hold across countries? Figure 5.6 shows that, overall, the same manufacturers are responsible for a high share of the infections in most of the top 20 countries with the most infections. We aggregate the data of the other 48 countries codes under ‘Others (48 CC)’ as well as the data of the rest of the 61 manufacturers that were not on the top 9 under ‘Others (61 Manuf)’. To calculate the ratio we divided the total number of infected devices in a country by the total number of infections attributed to a manufacturer. There is some variability, of course. In some countries, the share of ‘unknown’ is very high. Furthermore, in some countries the share of ‘others’ manufacturers is larger than that of the nine manufacturers. Still, in many countries the same manufacturers are present in the population of infected devices.

Table 5.2 quantifies this more clearly. We checked which manufacturers are present in the population of infected devices in each country. Meaning that the manufacturer at least appeared once in the data of that country. We can see that HikVision devices are in the infected population in 54 (79%) of the 68 countries in our measurements. Avtech devices show up in 47 (69%) of all countries. Those two together are present in most countries and




Figure 5.6: Share of manufacturers in top 20 countries with the most infections (countries ordered by number of infections)

they represent over half of all infections that we could attribute to a manufacturer.

This suggests that international collaboration among regulators in various countries is a feasible path. This would not only bundle scarce resources on the side of governments, but is also more likely to influence manufacturer behavior through collective action. An obvious starting point would be coordination at the level of the European Union. When we look at the distribution in the E.U. countries in Figure 5.7, we also observe the same nine manufacturers associated with most of the infections. We aggregate the data of the other 40 countries under the label ‘Others (40 CC)’ as well as the data of the rest of the 61 manufacturers that were not on the top 9 under ‘Others (61 Manuf)’. As before, to calculate the ratio we divided the total number of infected devices in a country by the total number of infections attributed to a manufacturer.

Table 5.2 also demonstrates that the locations of the manufacturers’ headquarters (HQ) are highly concentrated in China and Taiwan. This suggests another path for coordination, where the governments of those countries could help facilitate improved security practices in the manufacturing processes, in order to safeguard access to overseas markets thus this can give some leverage to governments to discuss with them their security postures since their IoT products are being imported to their countries.

In sum, the dataset gives a clear answer to our first two research questions. First, which manufacturers are associated with the compromised IoT across 68 countries? It turns out that just nine manufacturers are associated with about half of all infections. Second, how variable is the set of manufacturers across different countries? We find that—notwithstanding regional and country-level differences in consumer preferences, regulatory regimes, and market access—this pattern is remarkably stable across countries.




Figure 5.7: Share of manufacturers in E.U. countries (countries ordered by number of infections)

Manufacturer	HQ	Presence (%)	Share attributed infections (%)
HikVision	China	54 (79%)	28%
Avtech	Taiwan	47 (69%)	25%
MikroTik	Latvia	40 (59%)	7%
Xiong Mai	China	50 (74%)	7%
Synology	Taiwan	28 (41%)	3%
Merit Lilin	Taiwan	26 (38%)	3%
TP-Link	China	36 (53%)	3%
QNAP	Taiwan	34 (50%)	3%
Huawei	China	28 (41%)	3%

Table 5.2: Manufacturer presence across countries

5.5.2. DEVICES

Although our main focus is on device manufacturers within this study, it is informative to get a sense of which types of devices dominate the infected population. Figure 5.8 shows that, where we were able to ascertain the device type, almost 80% of the infected devices are Digital Video Records (DVRs) and IP cameras. As described in the method, as long as the devices were vulnerable to Mirai, they were considered infected. When checking the manuals, most of these devices had weak hard-code credentials. This is line with [166] work, which describe that guessable passwords vulnerable to attacks are used by the manufacturers in some of these device categories.

The Open Web Application Security Project (OWASP) describes default credentials as a top threat for IoT devices [215]. Mirai’s most famous attack vector is brute forcing




Figure 5.8: Top devices over the period of observation

attacks, and since Mirai’s source code was released attackers can easily add credentials to the code. Authentication in IoT devices sometimes is hard-coded or manufacturers use default credentials to set up a device for the first time, and this allows attackers to perform password guessing [61]. After the initial set up, most devices do not request to change these default credentials [200]. Therefore, manufacturers of these devices can help to fix most of the infections by implementing a better password management creating unique credentials per device.

5.6. UPDATES AND SECURITY ADVICE

Our third and final question is: What are manufacturers doing to remediate the security weaknesses of their devices? To answer this, we looked at the manufacturers’ websites and at manuals (see section 5.4). A wide-spread complaint is that many of the vulnerable devices never receive updates [137, 249]. We found that 37 (53%) of the 70 manufacturers present in our dataset had either firmware or software available to download. Of the 70 manufacturers, 30 (43%) describe a password changing procedure, and 18 (26%) have some security advice on how to make the device more secure and protect it from attacks.

The picture is a bit more positive for the top 20 manufacturers associated with most infections. Of these 20, 13 (65%) had some firmware or software available to download related to their devices, and 12 (60%) describe some password changing procedures, and 8 (40%) provide some advice to protect the device from attacks or make it more secure. The two dominant manufacturers, HikVision and Avtech, had firmware and software available to download in their websites.

Table 5.3 presents a summary of our collected data for the top 20 manufacturers. The FW/SW update column depicts whether or not a firmware or software was found available, the password changing procedure column shows whether or not we found a password changing procedure for the device, and the last column depicts if any advice to protect

the device was found or not (see more details in Figure 5.4). The data for the full set of manufacturers is provided in Appendix D.1. In sum, most manufacturers are making efforts to publish updates, password changing procedures, and provide security related advice. A significant group, however, does not provide one or more of these forms of support for protecting their devices. 81% lacks at least one of these three forms of support and 33% lack all three forms.

	FW/SW	Password changing procedure	Advice to protect the device
HikVision	Yes	Yes	Yes
Avtech	Yes	Yes	No*
MikroTik	Yes	Yes	Yes
Xiong Mai	Yes	No	No
Synology	Yes	Yes	Yes
Merit Lilin	No	Yes	No
TP-Link	Yes	No	Yes
QNAP	Yes	Yes	Yes
Huawei	Yes	Yes	No*
ZTE	No	No	No
Beijer Electronics	No	No	No
Zhejiang Dahua Technology Co., Ltd.	Yes	Yes	Yes
DrayTek	Yes	Yes	Yes
AVM GmbH	Yes	Yes	Yes
Domoticz	Yes	Yes	No
ASUS	Yes	No*	No*
Hichan Technology	No	No	No
ZKTeco	No	No	No
ZNDS	No	No	No
Sansco	No	Yes	No

Note: The asterisk in "No" means that multiple devices of this particular manufacturer were found in our data. For some of the devices the password changing procedure was found, but for others not. The same holds for advice to protect the device. See Appendix D.1 for more details for each device or category.

Table 5.3: Manufacturers offering software/firmware and security advice

Although these findings suggest broad manufacturer support for security, it is far from complete. It is often quite difficult for consumers to find and understand the relevant information. NIST's "Foundational Cybersecurity Activities for IoT Device Manufacturers" [115] emphasizes the importance of specific IoT product information to communicate to customers and how this communication is achieved. Information such as device support, lifespan expectations, end-of life periods, how to communicate suspected vulnerabilities during and after the life span a device to the manufacturer, security capabilities of the device or manufacturer services, how to maintain security after support of the manufacturer ends, type of software updates and whom will distribute them among others things are all examples of topics that could be communicated to users according to the NIST framework.

Moreover, NIST advice states that manufacturers should provide information on whether software or firmware updates will be available, when they will be available, and how customers can verify the source and content of the update (e.g. via cryptographic hash comparison).

To illustrate, on the Hikvision website [144], one can obtain the firmware of the device, but there is no explanation of how customers can verify the source and the content of the update. Similarly, Avtech's website [32], while providing firmware download options, does

not provide visitors with information on how to verify the authenticity of the firmware content either.

Although we were not assessing if manufacturer websites comply with the NIST framework, our brief examination of their content suggested that most do not offer all prescribed information, but a more systematic analysis is necessary to comprehensively assess all manufacturer websites.

During this analysis from an end-user's view, we also found that checking a manufacturer's website or manual is quite challenging in certain cases. Most websites focus on providing commercial information about devices, features, and comparison among devices. Finding manuals, support or updates might require numerous steps to achieve. In addition, the language used can be very technical. In a handful of cases, we also ran into situation where the products were discontinued by the manufacturer and we could only find the relevant device manuals on third-party websites. This pattern is aligned with the findings of [39]. Little security is provided by the manufacturers. All of this suggests room for improvement given that all these manufacturers are producing devices that are being compromised at scale.

5.7. RELATED WORK

CONSUMERS AND IoT SECURITY

An important area of IoT (in-)security research has focused on empowering consumers to consider the security and privacy implications of purchasing certain IoT devices. Vendors typically do not provide information on the security features and privacy sensitive characteristics of their products – information that may help consumers make more informed purchasing decisions – and when they do, it is often inadequate [37, 101]. Various studies have thus focused on developing security labels to better inform consumers [102, 192].

Privacy advocating organizations have also introduced valuable tools and guidelines to emphasize online safety and help consumers make more informed purchasing decisions, for instance see Mozilla's *Privacy not Included* guide [273].

The potential role of third parties in protecting consumers, for instance the role of ISPs in their capacity to mitigate IoT insecurity problems, at least as a short-term solution, has also been recently examined [24].

Nevertheless informative labels, consumer empowering tools, nor third parties like ISPs, can systematically prevent post-sale security issues in IoT products [24, 192]. They do not replace the necessity of engaging with manufacturers of (compromised) devices to get them to address the security problems of already sold or newly developed IoT devices.

REGULATIONS AND STANDARDS

Leverett *c.s.* [170] argue that existing sectoral regulators need to determine where IoT is present in their sector and to include them into existing safety and security regulations. They also highlight the need for transparency regarding products and vendors—to which our study is contributing. The European Union Cybersecurity Act provides a voluntary certification scheme for digital products, including IoT devices, in order to increase trust and security of these products [108]. Also, product liability could lead manufacturers to comply with minimum security standards in order to reduce their exposure [195].

These long-term solutions regulatory strategies, yet do not reduce the current influx of insecure devices, and our work presents an empirical approach to identify priority targets for governmental intervention.

INTERNAL MAPPING OF IIOT DEVICES

Several studies use internal network scans to identify IoT devices. One study [166] used the Avast Wifi Inspector to scan 16 million home networks and found 83 million connected IoT devices. To identify the manufacturers, the researchers matched part of the device MAC address with the public IEEE Organizationally Unique Identifier (OUI) list. Another study [301] created ‘IoT inspector’, a tool that users can run inside their home networks to label IoT devices and their manufacturers. Similar to [166] the authors use MAC addresses to validate vendors against the OUI database. A different approach was taken in [20], which fingerprints devices using information related to the Inter Arrival Time (IAT) of packets on the local network. This method was tested with just two devices in a lab setting.

All of these methods rely on user consent and privileged access to internal network data to identify manufacturers. This limits the scalability of the approach that is needed as a basis for governmental intervention, especially when representative measurements are needed across entire countries or markets. Therefore, in our study, we build on recent work on external mapping of IoT devices instead.

EXTERNAL MAPPING OF IIOT DEVICES

Numerous studies identify IoT devices in the wild based on external network scans. Most are based on developing fingerprints from known devices, *e.g.* in a lab setting, and then searching for these fingerprints in internet-wide scans. For example, one study builds fingerprints based on specific port configurations that are chosen by manufacturers [260]. The authors test their fingerprinting approach for 19 IoT devices and subsequently develop a hierarchical port scanning method to detect device types during external scans rather than probing whole port ranges. The approach assumes that end users will retain and not modify the specific port configurations of their devices used for fingerprinting. In [180], the authors fingerprinted

routers using the initial time to live (TTL) of two Internet Control Message Protocol (ICMP) messages to determine the brand of the routers' vendor. They highlight that the hardware distribution of different brands vary across Autonomous Systems. [218] proposed IoTFinder, which contains fingerprints for 53 devices that were developed from DNS traffic data and then compared these fingerprints to traffic from an ISP network. In [247], fingerprints are developed from a testbed setting, in this case for 96 devices belonging to 40 vendors. They then enriched their fingerprints with DNS queries, web certificates, and banners and detected IoT devices in an ISP and at an Internet Exchange Point (IXP). A different approach to generating fingerprints was presented by [121]. The authors searched the web for product descriptions of devices and then they automatically created fingerprints from these descriptions (e.g., rules to detect certain strings). This potentially scales better than generating fingerprints from analyzing the devices themselves or their firmware. However, [151] challenged the reproducibility of this method.

A common feature among these approaches is that they first develop fingerprints for a set of known devices under the control of the researchers and then conduct external scans with these fingerprints. Furthermore, some approaches—e.g., [218] and [247]—need access to ISP or IXP traffic in order to detect their fingerprints. This approach does not work for our problem of identifying a given population of devices in the wild, namely compromised devices. We cannot know which devices are in that population, let alone have them available in a lab setting for generating fingerprints.

Two other studies [130, 201] focused specifically on compromised devices. They identified the IP addresses of compromised IoT devices via attack traffic observed in darknet data. They did not develop fingerprints, however. The actual identification of the devices present at those IP addresses, was not conducted by the researchers. Instead, it relied on third-party data, most notably searching for the IP addresses in Shodan [258], a search engine that indexes a variety of internet-connected systems.

While we focus on consumer IoT, there is some overlap in approaches with the research on identifying industrial control systems (ICS) devices [92], which also relied on Shodan [258] and Censys [52]. Fingerprints were developed for individual ICS devices in order to track them over time, not for manufacturer identification. While also [296] developed a realtime ICS discovery system using ICSs protocols to discover ICS devices in the whole IPv4 space. They analyzed 17 ICSs protocols, and they did common requests that could fingerprint the ICSs devices based on the responses they obtained and that were unique to the protocols.

Like [130, 201], our study also uses attack traffic to detect the presence of compromised IoT devices, namely observing the Mirai fingerprint in darknet data. We base our analysis on

a longer data collection period of two months. For our device identification, we do not rely on a black-box third party solution like Shodan. This would make it impossible to explain to manufacturers via what method their devices were identified, nor gauge how accurate this method is. Explainability and accuracy—which includes knowing the method's inaccuracy—are key requirements for providing the government with the basis to select and engage manufacturers. Rather than relying on third-party services we develop our own fingerprints, as we explained in Section 5.4. Different from the other studies using fingerprints, we could not start with a set of known devices to develop the fingerprints. Rather, we need a method to identify manufacturers present in a given population of compromised devices.

5.8. LIMITATIONS AND FUTURE WORK

Our approach is to scan back an IP address from which Mirai scan traffic has originated moments earlier. A core assumption behind this approach is that the scan back will actually connect with the same device from which the attack traffic was observed. In reality, there will typically be multiple devices behind the same public IP address. Some, if not most, of those devices will not be publicly reachable. In theory we might be engaging with one of those reachable devices or with the router, either of which may or may not be the infected device. While we have no certainty that the device that we scanned back is actually the infected device, we have certain indicators that increase the confidence in our approach. First, attackers behind the Mirai infections recruit devices also by scanning IPv4 addresses for publicly-reachable devices, the same logic that we apply. So if they could infect a device, that device has to be visible in an active scan. In only 1.2% of the cases did we find fingerprints for different devices at the same IP address, consistent with the fact that in most cases only a single device was accessible from the open Internet. Second, the probability that we are scanning the Internet-facing router, rather than a Mirai-infected device behind the router, is severely mitigated in light of our data. Over 60% of the devices we identified were not routers. Where we did identify routers, these models were known to be vulnerable to Mirai. This brings us to a third indicator: all devices that we identified from the banners and Web-UIs were investigated to ascertain that they were actually reported to be vulnerable for Mirai. They were, without exception.

A second limitation is that our active scanning method did not use all protocols used by consumer IoT devices. It was limited to 'ftp', 'telnet', 'http', 'SSH'. This means we might miss devices that do not operate any of these protocols. Future research might expand the set of protocols and quantify what proportion we are missing as well as try to include all IoT related protocols to better understand the infections landscape.

A third limitation is related to the fact that we only scanned devices infected with a

variant from the Mirai malware family. Other malware families might bring into view additional devices. That being said, Mirai has been the dominant malware family for years and is still being detected as a leading malware family, responsible for 21% of the IoT infected devices [158]. Furthermore, it has been reported that the different IoT malware families often compete over the same devices [136], which suggests that the Mirai population is not systematically different from other families.

A fourth limitation is related to our use of the Mirai fingerprint to identify infected devices. There is an extremely small probability that this fingerprint occurs by accident ($\frac{1}{2^{32}}$, to be precise). That still leaves open the possibility that someone sends out this fingerprint on purpose. We are not aware of any use cases for doing this. Such an activity would not be part of a honeypot design for Mirai. In any case, it is unlikely that such technical corner cases would originate in substantial numbers from consumer broadband network (as opposed to research or hosting networks).

A fifth limitation impacts our assessment of the volume of infected devices in each country. IP address reassignment (a.k.a. DHCP churn) might impact the number of infections we observed per country. To minimize the impact of churn, we assumed that IP addresses are not reassigned multiple time per day. We count infections in 24hrs long sliding windows and with each batch of scans start a new count of unique IP addresses. This significantly reduces the risk of overcounting because of churn.

Another limitation is that we could not identify the manufacturer for a significant portion—roughly about half—of all infected devices. To the best of our knowledge, no other method for identify IoT devices in the wild has achieved better rates, but this is still a limitation of our work. As we discussed in subsection 5.5.1, the portion of unknown devices is unlikely to impact the pattern of concentration around nine manufacturers that we uncovered.

A final limitation is related to the fact that we are not sure when the websites or manuals of the manufacturers were updated. Hence, some of the security advice could have been recently added or not up to date. Moreover, we did not check if the firmware updates were actually solving the vulnerabilities of the device, but just if there was firmware or software available to download.

5.9. CONCLUSIONS AND DISCUSSION

The IoT ecosystem is complex and involves many different actors. Many observers have argued that the incentives in around IoT security are misaligned. [159, 295] There is a lack of adequate information available to consumers regarding the security of the devices that they are purchasing. The costs of security failures are often borne by other stakeholders than the owners of the device or the manufacturers. So there is a market failure here that

justifies government intervention. There is no single solution, of course. A recent step of the Dutch government has been a voluntary agreement with the main online electronics retailers to include in the product descriptions whether the product will receive security updates and, if so, for how long [211]. The current status is that many of these fields are still listed as 'unknown'. Many manufacturers are not supplying this information in their product description.

Any sensible strategy towards IoT security will have to change manufacturer behavior towards designing more secure devices. This is especially critical for the manufacturers associated with devices that have been getting compromised at scale in the wild. In this paper, we have investigated the manufacturers associated with the population of infected devices in 68 countries. We found that just nine manufacturers share about half of the infected devices across all countries. Notwithstanding the differences between countries in terms of consumer preferences, manufacturer presence in the market and regulatory regimes, this pattern also holds at the country level for most countries in the top 20 with most infections, as well as across European Union member states. Hence, policy makers can unite their efforts to target those to encourage them to improve their security postures. Most devices come, unsurprisingly, from China and Taiwan, the leading hardware manufacturers of the world. This concentration on the supply side of the market suggests that governments confronted with infected devices might engage their counterparts in China and Taiwan to change the behavior of the manufacturers in those countries, if only to safeguard their exports towards large markets in the U.S. and E.U.

Even though many manufacturers do provide security updates or advice, it seems that this is not enough to prevent and remediate the infections. This could be because of users' misaligned incentives [295], but it could also reflect that this support is hard to find and even harder to act on. The information on the support pages is fragmented. A user has to click different links, understand what files to download, and install them without a clear idea of what the new firmware version will or will not fix. Hence, there is room for improvement about what and how to present this information to users, as discussed in [115]. This would also reduce the cost that users have to incur to secure their devices.

The efforts that policy makers undertake can have an impact also outside their own jurisdiction. Think of how the E.U. became the *de facto* privacy regulator of the world, via the General Data Protection Regulation. Most websites adopted it globally, because it was more efficient than differentiating the setup for each jurisdiction [33]. If policy makers unify their efforts and the pattern of concentration on a handful of manufacturers holds, then a global impact is not unrealistic.

Retailers of IoT devices could also play a role, as countries such as The Netherlands are

proposing [211]. If users can return these devices to retailers, then these costs would lead the retailer to exert pressure further up the supply chain and create better security incentives for manufacturers.

Government involvement is currently underway. Many countries are introducing legislation or shoring up existing mechanisms to improve security. Our findings are a stepping stone for efforts by the Dutch government to engage the manufacturers found to be supplying most of the infected devices. Time will tell whether government pressure, in combination with empirical evidence of the problems caused by their products, is enough the start changing the security practices of these companies—and of the IoT market at large. These findings are based only on Mirai and we did not use all protocols used by consumers IoT devices, so future research could look into more IoT malware families and add additional protocols to have a more complete overview of the whole manufacturer landscape.

6

UNDERSTANDING PROTECTIVE DNS ADOPTION FACTORS

Protective DNS (PDNS) filters out DNS requests leading to harmful resources. PDNS is currently being promoted by various governments and industry players – some global public DNS providers offer it, as do some government-sponsored DNS resolvers. Yet, are end users even interested in adopting it? The extent of current PDNS usage, as well as the factors that encourage or discourage end-users' adoption, have not been studied. We found that overall PDNS adoption is minimal, though in some countries over 20% of the DNS queries are being answered by these types of resolvers. Four human-subjects studies were undertaken to understand end-user adoption factors: a survey with 295 consumers; 24 interviews with ISP customers offered a free PDNS after a malware infection; 12 interviews with public and private enterprise professionals, and 9 interviews with DNS technology specialists. We found that users are more likely to use PDNS if operated by their own ISP rather than the government. For enterprises, we uncovered that access to global threat intelligence, a layered security strategy, and compliance with regulations were the main factors for PDNS adoption. The DNS technical specialists highlighted broader challenges of PDNS adoption such as transparency and centralization.

6.1. INTRODUCTION

To access most Internet-connected services, the Domain Name System (DNS) is a crucial component [298]. The resolution of domain names to IP addresses has traditionally been provided by the recursive DNS resolvers of Internet Service Providers [230]. Since 2006, alternative recursive DNS resolvers, also called public DNS resolvers (a.k.a. open resolvers) have emerged [135, 230]. Companies such as Google, Cloudflare, Yandex and Cisco, and non-profits such as Quad9, have positioned their services as alternatives.

Every day, millions of new domain names are registered, some of which attackers use to redirect end users to harmful resources [109]. Hence, some of the public DNS resolvers offer services that aim to protect users by preventing the resolution of domains that lead to known malicious resources, like phishing or malware sites [4, 66, 67, 297]. This type of DNS filtering is called Protective DNS (PDNS) [197, 199].

Recently, some governments have started advocating that their citizens and enterprises should adopt PDNS as a security measure. The United Kingdom requires public sector organizations to use PDNS [197] and encourages adoption by private organizations. In the United States, the National Security Agency (NSA) and Cybersecurity and Infrastructure Security Agency (CISA) suggest that organizations use protective DNS as a best practice for their security strategy [199].

Some governments are actually backing specific PDNS services. Canada supports the CIRA Canadian Shield [65]. In January 2022, the European Commission announced plans to introduce DNS4EU [68], a recursive European DNS resolver service to protect citizens from malware, phishing, and other threats. The commission selected a public-private consortium to run the service and onboard 100 million users [292]. Australia has an initiative to encourage public sector entities that provide critical services to use AUPDNS, a DNS resolver that blocks cyber threats [275].

Although many governments and industry actors are pushing for PDNS services, their adoption critically depends on acceptance by consumers and enterprises. So far, solutions backed by a government have been imposed within government infrastructure itself, but not outside of it. The success of PDNS initiatives is based on voluntary adoption by users or their service providers. Service providers, in turn, have little incentive to adopt PDNS and provide it by default if there is no demand for it from their customers, since it also means losing a crucial data source for monitoring security threats on their network (e.g., observing customer DNS queries to botnet command-and-control servers). Even though adoption is critical to success, no research has examined whether citizens and enterprises have currently adopted PDNS and what factors drive or discourage adoption. Our research aims to fill this gap.

The closest related work asked home users about DNS over HTTPS (DoH) settings in Brave, Chrome, Edge, Firefox, Opera, and Android mobile operating system [208, 209]. These studies examined whether users changed their DNS settings after being told about encrypted DNS. DoH concerns user privacy, specifically which DNS provider can see queries. Our work explores user perceptions of DNS filtering by the resolver, which is previously unexplored and has implications for both the growing market in PDNS products and nation-level strategies for the uptake and utility of PDNS.

Our first research question is: what is the extent of adoption of public DNS and, in particular, PDNS? We estimate adoption by analyzing a dataset from the Asia Pacific Network Information Centre (APNIC) Labs[147]. From data from over 240 countries and territories on DNS recursive resolver usage and more than 15M average daily DNS queries over six months, we identify which portion went to PDNS providers. Earlier work underlined the importance of examining adoption factors for user-facing security and privacy technologies [9]. Hence, the main contribution of our study is focused on our second research question: what factors encourage or discourage the adoption of Protective DNS by users and organizations? For this question, we conducted four complementary human-subject studies. First, we carried out a survey in Prolific [225] with 295 participants to understand users' views on a PDNS service and what factors they would consider for adopting it. Where the survey captures the intention to adopt, we complement it with an interview study with data on actual adoption, with 24 customers of an Internet Service Provider who were offered to opt-in for a 'malware protection service' based on PDNS. Next, we interviewed 12 professionals in public and private organizations to understand factors to consider for adopting a PDNS resolver in an enterprise context. Finally, we interviewed 9 DNS technology experts who provided their reasoning on broader challenges before adopting PDNS resolvers. The main contributions of this paper are:

- By analyzing DNS recursive resolver usage of over 240 countries and territories and 15M average daily DNS queries, we determine the adoption of commercial PDNS resolvers in different regions (Asia, Africa, America, Oceania, and Europe) and countries.
- This research is the first to examine, across a variety of vantage points – namely users, ISP customers, enterprises, and experts – the factors that influence the adoption of PDNS as a security countermeasure.
- We find an adoption intention for 58% of users, but signal might overestimate demand, as only 9% of users signed up when the ISP we worked with offered them PDNS after they had suffered a malware infection.

6.2. METHODOLOGY

Given the recent push for PDNS, we first analyze a DNS resolution dataset for evidence on PDNS adoption. Next, we survey end users to learn about adoption factors. To complement the survey, we interviewed ISP customers who were offered a free PDNS service by their ISP. Then, we analyze interviews with professionals about why their organizations might adopt or not PDNS. Finally, we interview experts on the relevant factors for or against PDNS adoption.

6.2.1. RECURSIVE DNS RESOLVERS MEASUREMENT

APNIC dataset description. Asia Pacific Network Information Centre (APNIC) Labs performs a daily DNS resolver measurement to record users' sets of resolvers in DNS logs using a Google advertisement campaign [129, 134].

APNIC links the DNS resolver's Internet Protocol (IP) address and the user IP address to their autonomous system (AS) to identify the resolver type. If the resolver IP and user IP belong to the same AS, the resolver operator is most likely the user's ISP. This is counted as the resolver being in the 'same AS'. This omits public DNS resolvers. If the resolver IP address and the user IP address are in the same country, they increase the 'in country' count or 'out country' if the IP address of the operator of the resolver is not in the same country as the user IP. Public DNS resolver counts are excluded from 'in country' and 'out country' counts. Finally, if the resolver's IP address is associated with an AS of a public DNS resolver, they add the count to that resolver.

Processing and analysis of APNIC dataset. We employed the 'first use' resolvers — which is the first resolver seen for the user query in the DNS logs from January to June 2022 — to estimate which portion of DNS queries show PDNS usage. The data we obtained contains only the daily counts of DNS requests answered by each APNIC-labeled DNS resolver.

To determine if each public DNS resolver was a PDNS or not, we thoroughly examined their websites and service descriptions. We classified public DNS resolvers as 'Protective DNS (PDNS)' if they advertise themselves as protecting against botnets, malware, phishing, and spam. If not, they were categorized as 'No Protective DNS'. Cloudflare, Yandex, and Quad9 offer PDNS-enabled and PDNS-disabled services. Since we can only have access to counts, it is hard to determine which one the user is using, so we classify them as 'Possible Protective DNS'. Finally, we came across three cases — Free DNS, Level 3, and puntCAT — where we could not find information to determine whether they offer PDNS or not, so we categorized them as 'No information'. See Appendix E.1 for summary of the classification of the public DNS resolvers.

Next, the penetration of each DNS resolver category was computed as a percentage. We

divided the sum of the average daily unique queries per resolver type by the total average daily DNS queries. We calculated the percentage per country and aggregated different regions (Asia, Africa, America, Oceania, and Europe). China was undersampled compared to its Internet users, thus we removed it from our analysis.

6.2.2. PROLIFIC SURVEY

In August 2022, we ran a survey on Prolific [225] to gauge users' PDNS adoption intentions. The survey was created with Qualtrics. We paid proportional to the participants' completion time using the equivalent to the minimum wage from where the authors are based. PDNS was described to participants based on a literature review of the UK government's existing description [197] and the European Union tender for DNS4EU [68]. The survey design was informed by Fogg's behavior adoption model [123], which posits that motivation (M), ability (A), and trigger (T) (now 'prompt' [124]) affect how likely a behavior is to occur (in this case, opting in for PDNS).

Motivation focuses on the users' reasons for opting into PDNS. We asked participants about their 'perceived vulnerability' and 'perceived severity' to operationalize their motivations. Two questions concerned privacy, and a third examined the service's effectiveness against common threats. We also asked participants about what they regarded as significant threats, and whether they believed the service would be useful against them.

Users' skills determine the ability to perform a behavior. Time, effort, money, and pondering are elements of simplicity that increase ability [123]. We added questions about participants' ability to configure security on their devices, if they had other security methods, and use parental controls. Also, participants were asked if they would pay for the service. Finally, participants' awareness of comparable services was questioned, as a user may know about PDNS (i.e., have the ability), but not have the motivation to enable it.

Finally, according to Fogg's model [123, 124], triggers/prompts can be 'facilitators' that make the behavior easier, 'sparks' that inspire behavior, or 'signals' that remind the person to perform a behavior. A facilitator here could ease the adoption of PDNS or indicate its benefits – ISPs, DNS providers, and governments can facilitate DNS-blocking. Participants were asked which provider they preferred (Government, ISP, or commercial organization), but could add another.

Participants were also asked basic questions about their Internet usage and their computing devices. Open-ended questions let respondents explain their answers; some of these answers will be described in the results (See section 6.4). Demographic questions concluded the survey.

We added two attention-check questions, which all participants answered correctly. Before launching the survey, we ran two focus groups and a pilot (see Appendix E.2 for more

details). The survey included a measurement to record participants' resolvers' IP addresses to determine if they were using PDNS or not (see Appendix E.3 for more details). Appendix E.4 contains the whole survey protocol.

Participants. We calculated the number of survey participants using power analysis [122]. We included countries where more than 25 users were active in the last three months to calculate the total population. With a conservative estimate of 25% of the population proportion using PDNS and a 95% confidence level, 288 or more participants were required to answer the survey. Then we used a proportionate stratified sample of the same countries to collect our sample size. We collected data from 295 participants from 29 different countries. The participants' ages range from 18 to 66, with a mean age of 34. The stated genders of the respondents were 155 men, 135 women, and five who identified as another gender. 103 participants were located in America, 144 in Europe, 21 in Africa, 22 Asia, and 5 in Australia. The survey was only open to Prolific members who had approval rates of 90% or higher and had previously completed at least five studies.

Variables coding and ordinal logistic regression. All Fogg suggested variables namely motivation, ability, and trigger were included in an ordinal logistic regression model [161] (we performed two by two correlations, and the independent variables were not correlated). Stepwise, we extended the model to include additional variables – gender, age, and education – to produce a second, third, and fourth model respectively. In a final model, we added regions. We use the lowest Akaike information criterion (AIC) to choose the best model, considering that if an uninformative parameter does not explain enough variation, it should be removed [25]. We used as baseline the first model to compare the rest of the models. See Appendix E.9 for a summary of the variables of the final model with the lowest AIC value (model 2). The model aimed to predict which of these variables predict participants' PDNS adoption. The ordinal Likert scale responses to 'How likely are you to subscribe to Protective DNS if it were available today?' (slightly modified for home users willing to pay for the service) was the model's dependent variable.

We performed a factor analysis on Likert scale items measuring perceived vulnerability, perceived severity, and users' concerns (See Appendix E.4, questions 6–12, 24–26). All items loaded in their respective factor, so we computed their means for the regression model. The only continuous variable was perceived usefulness; the rest were categorical.

6.2.3. ISP CUSTOMERS INTERVIEWS

We partnered with a Dutch ISP from February to August 2021. 292 malware-infected clients were offered free PDNS by the ISP. Registering and consenting for the ISP service took about 2 minutes on the ISP website. 284 consumers received the invitation since 8 had email delivery difficulties. Of the 284, 259 (91%) did not enable the service, and 25 (9%) activated

it. After a month, consumers were asked to participate in a phone interview (See Appendix E.5 for the complete interview protocol). They were contacted via their subscription email address. 24 (8%) of 284 consumers consented to interviews. No compensation was offered for participating.

Nine of the 24 participants (37.5%) activated the service and 15 (62.5%) did not. Only 5 customers identified as female, the rest as men. The participants' ages ranged from 24 to 60 years old. We asked about their current security measures, how severe they perceive the possibility of someone abusing their Internet-connected devices, why they enabled or did not enable the service, and if they saw any drawbacks in using PDNS. The interviews were conducted in the participant's native language and recorded with their consent. The sessions lasted 10 minutes on average.

Two researchers independently coded the transcripts in Atlas.ti for thematic analysis [43]. The two coders utilized a sample of transcripts to generate initial codes for the themes. Informal discussions were used to ensure the reliability of findings as suggested by [45]. Over the course of twenty-three interviews saturation was reached (no new codes emerged from the interview). Five themes were found and they are presented in subsection 6.4.3. The frequency of each topic among participants is shown in Appendix E.8 along with code examples that helped group them into themes.

6.2.4. ENTERPRISE INTERVIEWS

Between April and July 2022, we conducted virtual meeting interviews with twelve professionals in charge of managing threats from malicious domains in enterprises. Except for one product manager, all were IT experts such as Chief Information Security Officers, security architects, and risk and IT security managers. The interviewees were not compensated for their time.

Via various social media accounts, we set out to recruit participants. Out of the total number of participants, 5 were employed by the government, 2 by banks, 2 by universities, 1 by a cable business, and 2 by Internet Service Providers (ISPs). One of the enterprises has its primary operation in America, one in Asia, and two have global operations; the rest were based in Europe. One of the ISPs serves 30,000 customers and the other 2 million. The rest of the enterprises are in charge of managing somewhere between 200 and 40,000 endpoints.

The interview questions asked if the practitioners were aware of PDNS, the pros and cons of implementing this security solution in their enterprises, and the factors to consider in using it. If necessary, participants were provided a definition of PDNS. Appendix E.6 contains the interview protocol.

For the interviews, English was the language of choice. The recordings lasted 43 minutes on average. Recordings were transcribed and anonymized, leaving out participants' names

and affiliations. Two researchers independently coded the transcripts in Atlas.ti for thematic analysis [43]. The two coders utilized a sample of transcripts to generate initial codes for the themes. Informal discussions were used to ensure the reliability of findings as suggested by [45]. Over the course of seven interviews saturation was reached (no new codes emerged from the interviews). Five themes highlighted the key subjects discussed by enterprise participants, and they are presented in subsection 6.4.4. The frequency of each topic among participants is shown in Appendix E.8 along with code examples that helped group them into themes.

6.2.5. EXPERTS INTERVIEWS

Nine DNS technology experts participated in semi-structured virtual meetings interviews (one interview was in person) over the period between March and May 2022. The interviewees were not compensated in any form.

From the RIPE DNS working group's open mailing list, we collected the email address of 28 DNS specialists debating DNS4EU. Nine agreed to the interview. DNS experts were from a range of countries, including the European Union. They have a variety of DNS-related experience, including building open source DNS resolver software, working with country-code top-level domain registries and participating in the development of Request for Comments (RFCs) related to DNS. Five experts describe they had global experience in DNS and four at the European Union level.

The interview questions focused mostly on learning what the DNS technology experts describe as 'PDNS', their opinions on this security countermeasure, how PDNS differs from other security countermeasures, its benefits and drawbacks, and their thoughts on governments' initiatives (see Appendix E.7 for the complete interview protocol).

For the interviews, English was the language of choice. The recordings lasted 51 minutes on average, and transcripts were anonymized. Two researchers independently coded the transcripts in Atlas.ti for thematic analysis [43]. The two coders utilized a sample of transcripts to generate initial codes for the themes. Informal discussions were used to ensure the reliability of findings as suggested by [45]. Over the course of six interviews saturation was reached (no new codes emerged from the interviews). Eight themes highlighted the key subjects experts discussed, as presented in subsection 6.4.5. The frequency of topics among participants is shown in Appendix E.8, along with example codes for how they were grouped into themes.

6.3. ETHICS

The protocol of this research was approved by the human research ethics committee of our institution (Reference number: 1920). Prolific participants provided their consent to

Table 6.1: DNS Resolvers Usage (Period: January to June 2022)

Region	avg daily queries	Non Public DNS resolvers			Public DNS resolvers				Total
		% Same AS	% In country	% Out country	% PDNS	% Possible PDNS	% No PDNS	% No Info	
Africa	1,671,192	58.2%	9.3%	1.2%	2.0%	2.0%	26.0%	1.3%	100%
Oceania	73,443	83.0%	5.3%	1.3%	1.0%	2.3%	7.0%	0.1%	100%
America	2,804,980	65.0%	9.2%	1.3%	0.9%	3.1%	20.2%	0.3%	100%
Europe	1,758,927	75.2%	7.6%	1.0%	0.9%	3.2%	12.0%	0.1%	100%
Asia	9,023,027	59.0%	20.0%	1.0%	0.8%	2.0%	17.0%	0.2%	100%

Note: % **Same AS**: Percentage of average daily queries which resolvers ARE in the same AS as the users and NOT known public DNS resolvers. % **In country**: Percentage of average daily queries which resolvers are NOT known public DNS resolvers and NOT in the same AS as the users, but ARE geolocated in the same country as the user. % **Out country**: Percentage of average daily queries in which resolvers are NOT known public DNS resolvers and NOT in the same AS as the users but, and NOT geolocated in the same country as the user. % **Public DNS resolvers**: percentage of average daily queries which are answered by resolvers as categorized in Appendix E.1.

participate in the survey. We informed them that we were collecting their IP addresses and their DNS provider (who responds to DNS queries). Participants were reminded that they could stop at any time. In exchange for the time spent completing the survey, we paid proportional to the participants' completion time using the equivalent to the minimum wage from where the authors are based. All interviewees in the three interview studies gave consent for the interviews and recording. They were reminded that they could stop the interview at any time. For the ISP customer interviews, customers' personal information never left the ISP's premises. In accordance with the terms of service and in agreement with the ISP's privacy team, we obtained an anonymized dataset for our data analysis.

APNIC collects users' DNS resolvers through advertisements. Users arriving at these websites have agreed to their terms and conditions, including the use of adverts. The APNIC data we obtained was anonymized and cannot be traced to individuals.

6.4. FINDINGS

6.4.1. PROTECTIVE DNS ADOPTION

We analyze what percentage of average daily DNS queries is answered by different DNS resolvers. In Table 6.1, we split non-public and public DNS queries, and we present the regional distribution. PDNS resolver adoption is low in all regions compared to resolvers in the user's Autonomous System (AS), which are normally operated by their ISP. We note that Africa is the region that uses PDNS the most, with 2% of all queries answered by this type of resolver. Oceania follows with 1%, America and Europe with 0.9% each, and Asia last with 0.8%. In all regions, OpenDNS is the most popular PDNS resolver.

When looking at the country level, a small number of countries – with Israel (IL) in the lead – have more than 20% of the average daily requests answered by a PDNS resolver (see Appendix E.10). The APNIC data does, of course, not show who took the action to set up PDNS: the end users themselves or others, such as the network provider. This high adoption in some countries does underline the importance of comprehending end-user opinions about

this service, since so many queries are already routed through them.

6.4.2. PROLIFIC SURVEY

Out of 295 participants, 13 were extremely unlikely to use the PDNS service presented to them, 34 were somewhat unlikely, 76 were neither likely nor unlikely, 120 were somewhat likely, and 52 were extremely likely to use the service. When the last two groups are added, the intention to adopt is 58%.

Participants' PDNS Awareness, first impression, and comparison to similar services.

During the survey, participants were introduced to a 'Protective DNS service' (See appendix E.4) that safeguards their devices against untrustworthy websites and malicious software which uses DNS to perform this task; we balance this explanation with clarification that, for instance, the service cannot stop all threats.

Participants were asked if they had heard of a similar service to 'Protective DNS service' before filling out the rest of the survey. 102 (34,6%) participants said they had heard of similar services, 121 (41%) did not hear about it, and 72 (24,4%) participants were unsure.

Participants were asked to describe the service using adjectives from a list; they could add more. Examining the responses of the 121 participants who had never heard of a service like this, 71 of these participants said the service would be helpful; 66 said it would be useful; 52 said it would be secure; 21 said it would be easy to use; 24 said it was confusing. 7 out of 121 participants thought the service was unnecessary for them, while 2 thought it was unclear. The service was viewed as useless by 1 participant. Anxiety was one participant's first reaction to describing the service.

Additionally, we looked at whether these participants perceive the PDNS service useful for the top security threats they face. Among the 121 participants that stated having never heard of a service like this, 59 (49%) considered the service very useful, and 27 (22%) participants considered it extremely useful.

We followed up with the 102 participants who had heard of similar services by asking if they were using a service similar to the 'Protective DNS service' described. Out of the 102 participants, 29 acknowledged using a comparable service. Therefore, we asked what the service's name was. Compared to the presented service, 11 out of 29 participants cited other public DNS resolvers. Four participants mentioned Cloudflare, two OpenDNS, one Adguard secure, one Comodo secure, one DNS filter, one Next DNS, and one mentioned a DNS filter for ads on their phone. Fourteen out of these 29 participants compared the service to other security countermeasures, such as antivirus, Internet browsers such as Chrome and Opera, cloud security providers like Sophos, and using Pi-hole. Remaining participants could not recall the name of the service they were using.

We further investigated the 11 participants who claimed to use public DNS resolvers

against our DNS data. All participants' DNS resolvers were in broadband ISPs in the same AS as where the participants' IPs were located, except for one participant who mentioned using OpenDNS but was actually using Yandex. We did not gather DNS measurements for the two participants who mentioned using Cloudflare and one participant who mentioned using Comodo Secure. These could be possibly explained by different factors: two respondents used their phones (which may not use the PDNS service), and respondents may have taken the survey from a different network than where they set up PDNS. We also cannot rule out socially-desirable answers [265], but mentioning DNS providers implies respondents were aware of the service.

DNS measurement results. As described in subsection 6.2.2, we included a measurement to capture the DNS resolvers of participants. We collected measurements from 285 of the 295 participants. Of those 285 participants, 208 (73%) had resolvers' IP addresses from broadband ISPs in the same AS as the participants' IPs. 28 (10%) participants had resolvers in a different AS than their IP address, but that AS belonged to the same ISP as the participants' AS. 22 (8%) participants had Google as DNS resolver. 19 (7%) participants out of the 285 were using a PDNS or possible PDNS resolver. Remaining participants used security vendors such as Akamai or Fortinet.

PDNS resolver users have never heard of or used a similar service, with one exception. This could mean another household member set up the service or that their network provider is re-routing DNS requests to PDNS resolvers. These findings highlight the need of measuring a phenomenon rather than depending merely on participant responses.

Who should provide PDNS? We checked all participants' responses about which provider they wanted for 'Protective DNS'. 156 (53%) participants chose their ISP, 100 (34%) a commercial provider, 24 (8%) their government, and 15 (5%) participants chose others (referring to non-profits and independent organizations that focused on privacy).

Internet Service Provider. We looked into why 156 participants chose their ISP to provide this service. 49 of 156 participants stated that ISPs were the most logical provider, as ISPs have the most understanding of current threats, can benefit from enhanced network security, and already provide their Internet connection. 37 out of these 156 participants mentioned trust. Twenty-three of 156 participants described that they have a contract with the ISP and that such a service can be bundled with it. 18 other participants who selected their ISP provided privacy-related justifications, with one remarking '*They can already access my internet history, so it would make no difference to my privacy*'.

Another group of 11 participants said they chose their ISPs due to their role in Internet connectivity. Eight participants were unsure or simply opposed to another option, so they did not consider a different party. Seven participants said the ISP's proximity and

ease of communication made this party appealing. Three respondents mentioned PDNS's affordability if their ISPs implemented it.

Commercial company We looked into why 100 participants chose a commercial provider. 41 participants mentioned that they would trust a commercial company more than their ISP or government. A participant stated *'I don't trust the government. A commercial company will be more transparent in what it does than the government or my internet provider'*. Twenty-three of the 100 participants also expressed that commercial companies would have better know-how, better resources, and staff than their ISPs or governments to carry out this task. Another group of eight participants mentioned reasons related to the incentive of profit that commercial companies have, with this leading to better service, as highlighted by one participant, *'There is room for competition among companies, you have options to change if the service doesn't meet expectations or doesn't align with what you think is important when it comes to your data, privacy or safety online'*.

The remaining participant's reasons to choose a commercial included simply preferring a commercial company (6 participants), having a good experience dealing with a commercial company rather than their government or ISP (5), they handle personal data better (4), the service will be cost-efficient (3), no other choice was appealing (2), easier to switch (2), commercial companies care about their reputation, so the service will be good (2), and the rest of participants did not state a reason for their choice.

Government Only 24 participants chose the government as their preferred provider. Fourteen of these 24 participants stated that they trust their government. Four participants stated reasons related to the government not having any economic incentive to offer the service, so they would not use their data to make a profit. A participant stated *'I feel the others [ISPs and commercial companies] would focus on making money - [I] feel like the government wouldn't use it [PDNS] as profit-making scheme'*. Of the remaining participants, reasons were raised that the government should be responsible for protecting its citizens (2 participants); government has access to their data (so it would not matter to them if the government offers the service) (2), and; government would have more resources for providing PDNS (1).

Least preferred provider We questioned participants about which provider they would not choose for Protective DNS. One hundred and ninety-one (65%) would not choose their government, 51 (17%) would not choose a commercial company, 31 (11%) participants would not choose other parties (mainly companies they do not know or trust and with low reputation), and 20 (7%) would not choose their ISPs.

Participants provided different explanations of why the government was their least preferred provider. Ninety-six out of the 191 participants mentioned distrust. Another group of 34 participants, expressed privacy concerns. As mentioned by one participant *'It would be*

like *Big Brother watching you*'. Nineteen participants expressed concerns about the ability of the government to deliver a quality, efficient or effective service. Fifteen participants mentioned that they would not choose the government because they might use this service for censorship. A participant said '*No government should have full access to everything their citizens do online. I don't want to be in a bootleg China/North Korea/Russia*'. Other groups mentioned they could block websites for political or personal reasons. Some said this gives the government too much power.

Twelve participants said this was not the government's role. One participant stated '*Doesn't seem like it's something in their wheelhouse*'. Two participants said it would be hard to complain if the service went wrong since the government is hard to reach. The other participants did not explain their choice.

Explaining adoption. We used ordinal logistic regression to predict PDNS adoption. As mentioned in subsection 6.2.2, we incorporated all Fogg's model variables and evaluated models with gender, age, education, and regions. These models have no significant variables, except for the one including gender. The model including gender as a control variable has a slightly lower AIC value (800.70) than the one with Fogg's factors only (800.85). Thus, we report that model (see section E.11 for a summary of the significant predictor variables).

Concern, perceived severity, self-installing security in devices, use of parental control, and PDNS awareness were not significant. To understand the significant model coefficients intuitively, we calculated the odds ratio [269].

As perceived vulnerability to malicious software and data theft increases (β 0.234, *OR* 1.264, p 0.1), the odds of adopting Protective DNS increase 1.264 times. Fear, to use Fogg's terminology [123], could motivate adopting PDNS to avoid threats.

As the perceived usefulness of the service to address threats participants recognize as important to them increases, the likelihood of adopting a PDNS service increases by 2.539 times (β 0.932, *OR* 2.539, p 0.1). This suggests that users who perceive value in PDNS are more likely to use it.

For participants who already have security measures (Antivirus, Firewall, Ad blockers) (β 0.774, *OR* 2.167, p 0.1), the odds of adopting PDNS are 2.167 times higher than for those without any security measures.

For participants willing to pay for the PDNS service (β -0.545, *OR* 0.567, p 0.1), the odds of adopting it are 0.567 times lower than for those who did not want to pay. This outcome was somewhat unexpected, so we looked at qualitative responses. One hundred twenty-two (41%) were willing to pay, whereas 173 (59%) were not. Out of 122 people willing to pay, 103 (85%) considered the service 'very useful' or 'extremely useful'. Why participants wanted to pay but indicated they did not want to use the service is puzzling.

This could imply, as Fogg [123] suggests, that cost in itself can reduce the Ability to adopt a new behavior (even if it does not completely diminish the possibility of change in behavior).

For participants in the group who chose their ISP as the preferred provider (β 1.091, *OR* 2.976, p 0.1) the odds of being more likely to adopt PDNS increased 2.976 times compared to participants who chose the government. For participants who chose a commercial company as the preferred provider (β 0.965, *OR* 2.626, p 0.1) the odds of being more likely to adopt PDNS increased 2.626 times compared to participants who chose the government. Interestingly, the ISPs and commercial companies acting as ‘facilitators’ rather than the government had a bigger effect size to predict the likelihood of PDNS adoption.

The control variable gender was significant. Females were more likely to adopt PDNS than men. As our instrument was not meant to measure gender differences and this was a control variable [146], we refrain from strong claims.

Additional information to decide to opt-in to the service. Participants were asked what information would help them subscribe to PDNS. Privacy policies and data use were popular subjects. Another topic was the service’s cost and effectiveness against intended threats. Participants also mentioned cancellation policies, expert reviews, reports on what the service protects, whether the service can be turned off, whether the service affects network speed or device operation, general terms and conditions of the service, customer reviews of the service, reputation and trustworthiness of the provider, how simple it is to use, and why the service is needed in addition to other security measures. Additionally, the time to set up the service may be significant to communicate to users, since 212 (71.9%) participants were willing to devote only one to twenty minutes to subscribe.

All in all, from the Prolific survey, we learned that the motivation elements for users to adopt the service were the perceived vulnerability and perceived usefulness of the service. From the ability construct, the cost was a significant factor. However, who is the provider of the service, the trigger/prompt [123, 124], plays the most important role.

6.4.3. ISP CUSTOMERS INTERVIEWS

Unlike the Prolific survey, we interviewed 24 ISP customers who were offered PDNS and suffered a malware infection three months prior to the interview. Nine of the customers we interviewed adopted the service provided by the ISP, while 15 did not.

According to Fogg’s concepts fear of something bad happening can act as a motivator to enact a behavior [123], in this case, the perceived vulnerability (malware infection) could motivate choosing the ISP PDNS. In the Prolific survey we observed that as perceived vulnerability to malicious software and data theft increased, the likelihood of adoption increased. In reality, this did not happen. Surprisingly, just 25 (9%) of 284 customers opt-in to the service. This may be because the survey measures intention, not actual behavior,

benefits were not communicated clearly, or users did not want the service even after suffering a malware infection.

Concerns about the service. Twenty-one customers voiced concerns about the service. 16 customers worried about privacy or data use. *C15 expressed 'Naturally it [PDNS] might have some implications for your privacy. I think that if you want to be sufficiently protected that necessarily comes at some cost to your privacy. I think that is something you have to take for granted; it is something that is inevitably linked'.* Two consumers expressed concern over their lack of understanding of how the service operates. Two participants were concerned about the service's effectiveness and the ISP's responsiveness in the event of a problem.

Reasons for non-adoption. 15 customers declined to use the ISP PDNS for different reasons. One person opted to control his own security, and another did not need the service. Another participant said that he was using work-provided equipment, so he did not enable it. Since the service requires no installation on any device, the customer may not have understood how the service operates.

Four users stated that they already had other software installed, namely antivirus, that protected their machines. To illustrate *C19 said 'Well, I have an antivirus program on my laptop that then stops everything that comes in from viruses, I think that is enough actually'.*

Three participants reported that they tried to follow the instructions but were unsuccessful in turning on the service. Additional justifications from two customers were that they didn't understand how the service operated.

A customer mentioned the possibility that the service may prevent accessing something he needed. Another client expressed that he did not enable it because the service could be billed later, and one customer forgot to activate the service.

Reasons for adoption. The nine consumers who enabled the service also provided a range of justifications for doing so. According to two customers, they enabled the service on their ISP's recommendation. One of these participants also expressed his fear of viruses. One customer stated that he considers the service useful as long as there is no payment involved. Another customer said that he thought an ISP could adopt security measures more quickly than an individual customer could.

One customer, *C1* said that he enabled the service because *'[it] automatically protects all devices that are connected to your router, that saves a lot of hassle'.* Three customers stated that they enabled it for a 'feeling of safety'. One client claimed that he enable the service to prevent malware from spreading.

Beliefs on abuse. Fifteen of 24 customers discussed various consequences of the misuse of their devices. Given that all participants have suffered a malware infection prior to the

interview, it seems that not all customers perceived it as a major event. Five out of these 15 customers enable the service. Two customers worry about data theft. One user claimed that his devices could spread malware, another said identity fraud could occur, and another said his network could be made accessible to the public.

Ten consumers who chose not to use the service still thought that device misuse would have consequences. Data theft was cited as the primary impact by six consumers. The remaining customers mentioned phishing, viruses, and the rise of hacking as consequences. Even though they believe malicious software is dangerous, they did not activate the service. The ISP sent a message inviting them to use the service for free but failed. The ‘spark’ [123] for motivating behavior may not have resonated with all of the customers.

Trust. Trust was also a topic mentioned by 11 customers. Six users that enabled the service trusted their ISP to do this job. Three individuals who did not enable the service expressed trust in the ISP, but they were the ones that tried to enable the service but failed. One participant believed that the ISP was a reputable party, but still chose not to enable the service since they had other measures in place. Conversely, one participant expressed distrust in the ISP.

6

Apart from ‘perceived vulnerability’ (which some customers cited as a reason to adopt the service and this study demonstrates that a smaller proportion of participants actually adopts the free ISP PDNS service), and ‘use of other security countermeasures’ (which actually lead participants to not adopt the service and perhaps think the service was not useful), these results support the findings of our survey.

Customers adopted the ISP PDNS because their ISP offered it, showing how important the trigger was. The service’s perceived usefulness also drove adoption. On the other hand, participants who did not use the service did not consider it useful for their circumstances (e.g. I have a device that is managed by my employer). Also, fear of future costs was listed as a reason for not opting in. Participants’ concerns correspond with the survey, being privacy a predominant topic. According to the instructions, some participants attempted to enable the service, but they were unsuccessful. Hence, there is space for improvement since as Fogg [123] suggests, effort and time spent can influence the Ability to conduct a behavior.

6.4.4. ENTERPRISE INTERVIEWS

We interviewed twelve professionals, as described in subsection 6.2.4. Two practitioners acknowledged that their organizations made use of a PDNS resolver. In addition, *P9* admitted that his organization had used a PDNS resolver, but no longer does because of the cost.

Two participants stated that their organizations run their own DNS resolvers and that they were filtering domain names at that level. One mentioned filtering Domain Generation Algorithms and the other a list of malicious domains. This shows that some enterprises can

deploy PDNS internally, and they also mentioned that would not use an external provider.

Two practitioners were unaware of PDNS, while the other participants knew about PDNS but did not use them for a variety of reasons. *P4* stated that their organization values do not align with this measure, *P6* said cost was an issue, and *P7* mentioned that they consider that DNS filtering is not always a viable security countermeasure. The two participants who work for Internet Service Providers were implementing services that offer DNS protection as an opt-in service.

Reasons to implement PDNS. The two participants who confirmed utilizing PDNS, mentioned the global threat intelligence as justification. The service gathers data from many businesses around the world, offers visibility on attacks, and prevents them.

P8 said that as no security solution is perfect, they added this extra layer of security. *P10* said that they chose PDNS service because it was straightforward to implement globally and because safeguarding their reputation was vital *‘We have a big name provider that it is in charge of filtering our DNS queries outside the organization. We use it because of the capability of the provider to deliver the service around the world since we have a lot of countries... it is not only cost.. for all organizations cost is important... Also, our reputation is important, so we take all the security measures that are possible’*. Also, *P9* indicated that their organization used PDNS in the past due to the value of global threat intelligence.

There were two organizations that added filtering to their own DNS resolvers, *P5* and *P1* cited having an in-depth defense strategy as the primary justification. *P1* stated, *‘we do filtering because of an in-depth strategy of protecting different layers... I think our organization and my colleagues tend to gravitate to just blocking stuffs’*. *P1*, however, claimed that the blocking that is now occurring in their resolver was out of date and was done with a static list.

P11 described that the main reason for offering this service to its ISP customers was that they believed that security was important. *P12*, on the other hand, said that the ISP implemented PDNS because its government mandated to block Child Sexual Abuse Material (CSAM) and they had to comply quickly. Since the solution was already in place, they saw the opportunity to offer businesses other types of blocking.

Factors to consider for adoption. We questioned participants who weren’t currently utilizing PDNS about what they would consider before implementing the service. Many factors were mentioned.

P9, claimed that cost and service efficacy were the two most important factors to take into account. He said that the fact that they ceased using the service was due to their inability to afford this security countermeasure. Two other participants, also identified cost as the primary determinant. *P6* said they depend on public funding to invest in their security

infrastructure. The same participant noted the need to examine this solution relative to their existing infrastructure.

P2 said that to evaluate the service's added value, they must consider its effectiveness. *P7*, on the other hand, who thinks DNS blocking might not always be a viable security countermeasure, said efficiency was the most crucial consideration. *P3* mentioned they would consider the organization's threat model and red teaming advice.

P4 stated that organization's values must be considered. When openness and transparency are desired, it may not be good to restrict domain names for the staff. Although not questioned, *P11* mentioned that the main consideration in adopting PDNS for the ISP where he works was consent. They had to consider all legal factors and build a way to obtain customers' consent to provide the service.

Concerns. Participants from adopters and non-adopters organizations raised different concerns, thus we separated them from adoption factors (See subsection 6.2.4).

Five participants expressed concern about false positives. However, none of the participants who were using a third-party PDNS said that a false positive had caused a disruption in their daily operations. *P10* stated, *'We experience it [false positives], but no frequently, but there were some hits. They have mechanisms to report it and the provider has excellent SLAs and we have ways where we can just make changes'*.

Additionally, *P3*, expressed concern about the time they would need to spend troubleshooting false positives. The trust a company places in a third party to manipulate the DNS responses was also brought up by *P5*. The service's transparency on what is being blocked was the main concern of *P6*. *P1* also stated that privacy was an issue since they are a privacy-conscious organization. *'When talking about blocklisting there are some concerns... because we have a lot of employees and all their traffic is passing within our network... even when they are working at home... As there is security consciousness, there is also very much a privacy consciousness on the end of our users...'*

Because DNS blocking may not be successful in all circumstances, *P7* expressed that his main worry would be that the organization would experience a false sense of security.

P11 was concerned that the service only protected devices linked to the ISP's router. If customers' phones are connected to a separate provider, they may get infected and customers might doubt the service. Second, the participant noted customers may not know the added value of the service because it does not provide reports of what is being blocked.

Government PDNS. Participants were asked if they would adopt a government PDNS. According to *P8* and *P10*, commercial PDNS services are global, while government initiatives are country or region-specific. For instance, DNS4EU will cover Europe. Commercial PDNS solutions provide them with improved threat coverage as a result. Both participants

said their organizations would use PDNS if required by the government, as they comply with other regulations.

P1 and *P5* indicated that it would be preferable if the government shared block lists that businesses could use on their own. *P1* highlighted ‘*If it would be a list that I could implement myself, then I would be interested . . . because then you can just also weed out filters that you may find too intrusive . . . and [have] more control of the actual blocking taking place*’.

P9, whose organization discontinued using a PDNS resolver for financial reasons, believed that these initiatives are a good concept and that they would explore adopting them. Both *P3* and *P6* agreed with *P9* that these projects are beneficial, and *P6* added that they are beneficial as long as organizations are free to set them up any way they see fit.

It comes down to who consumers trust, according to *P11*, and personally, he would put more faith in his ISP than government PDNS initiatives. *P12*, on the other hand, stated that he had contradictory opinions; on one hand, he dislikes government PDNS initiatives, yet the more security the better.

The majority of enterprise participants stated that their organizations employ PDNS because of the additional layer of security and global threat intelligence it provides. In addition, several other considerations for PDNS adoption were brought up, including PDNS efficiency, the organization’s threat model, the organization’s values, and cost. Consistent with the Prolific survey, these mentioned factors suggest that some sort of perceived usefulness depending on the characteristics of the organization as well as perceived vulnerability (threat model) might play an important role in PDNS adoption for enterprises. The Prolific survey results demonstrated that the cost reduced the likelihood of PDNS adoption (as with the Fogg model [123]). We have evidence of one enterprise stopping using PDNS due to the inability to pay for it. Across our studies, this indicates that the cost of the service is a concern for individuals and enterprise customers alike. Some of the enterprise participants’ concerns were transparency and privacy even though they might be adopting PDNS on behalf of their users. These concerns overlap with the concerns expressed by ISP customers. Trust in the provider was also mentioned by one enterprise’s participant as an important factor. This topic was discussed by ISP customers as well as it stood out in the survey as one of the reasons to decide to opt for a certain provider of the service.

6.4.5. EXPERTS INTERVIEWS

Factors for adoption. For users to adopt PDNS, *E1*, *E2*, *E3*, *E4* and *E7* emphasized the importance of awareness. One expert said DNS is beyond the understanding ‘common Internet users’. *E4* said that users usually stick with the default DNS settings offered to them, and it is hard to educate them on changing those settings. Thus, how easy it is to set up a PDNS resolver may affect its adoption. As part of awareness, experts described that it is

critical that users grasp PDNS policies, what they are signing up for, and what is blocked, what they are protected against, and how their data is used. E2 highlighted *They [users] should check who is providing it [PDNS]. . . will be an entity they trust? . . . if there is filtering what are the policies to turn it on and off . . . in general what users do is just to buy security . . . someone is selling a security tool, they turn it on, and then they forget about this. . . so this is, unfortunately, the average degree of awareness*.

Most enterprises prohibit access to particular internet resources using next-generation firewalls or proxy servers, according to E1, E2, E4 and E6. E4, E6 and E9 emphasized that the organization's size, security strategy, or network requirements may drive the adoption of PDNS. *'The level of filtering that can take place in DNS and especially in enterprise environments...well depends on the jurisdiction, will depend on the nature and strategy of the organization, size of the organization, will depend on the way that they're patronizing their employees or trusting them . . .'* E4 said. E1, E7, E9 mentioned it is vital to know where an organization's DNS data goes when employing third-party resolvers. Service level agreements, according to two experts, are crucial because the functioning of the organization will depend on an outside party. E5 suggested that enterprises should consider performance, ease of deployment, and maintenance when implementing a PDNS service.

The key issue raised by experts in regard to ISPs is their lack of incentives as they have nothing to gain from DNS blocking. ISPs may have DNS systems that enable PDNS, however, filtering in resolvers brings maintenance costs and no revenue. E3 suggested ISPs might adopt DNS filtering to offer to their customers if they could generate money by protecting users from DNS abuse.

According to five experts the main reason why governments should provide PDNS services to society is that doing so is in the public interest and for the benefit of society. E7 noted that some governments are interested in supplying PDNS, citing the DNS4EU initiative, as an attempt to dispel the notion that important infrastructure like DNS is run by unrelated commercial organizations with distinct objectives, and not adhering to the same European Union regulations.

Provider. PDNS is offered by numerous commercial public DNS resolvers. The majority of experts, however, concurred that the government should play some role. As it is in the public interest to prevent DNS misuse, the government is the appropriate party to provide DNS alternatives. They do not have a corporate reason to protect DNS requests above the interests of society. However, E6 questioned if the government could compete with private companies. E2, E3 suggested that ISPs should provide PDNS because most customers' Internet connections go through them. According to E1, any private organization can provide it. E9 recommended a federated effort, so no single entity would control DNS queries.

Limitations of PDNS. Six experts agree that PDNS's main drawback is that it is not a perfect solution and 'will not catch it all'. E1 stated, *'If you were to rely solely on DNS base security solution, you are going to run into problems because not everything will rely on DNS lookups in order to get the payload in, and if you will assume that you are protected, then you are not'*. Due to the dynamic nature of DNS, where attackers may use domain names briefly before a PDNS provider loses sight of them, the solution's success will depend on how accurate the threat intelligence is, according to E6. Experts also warned against using PDNS as their single security measure.

Another drawback of PDNS, according to E1, is false positives. E7 stated that users could get around using standard methods like virtual private networks.

Types of blocking. Despite the limitations of PDNS services, preventing abuse was mentioned by five experts. Most experts agreed that restricting domains for security is an unambiguous strategy. However, E1, E4 raised that a resolver could block categories based on keywords that could lead to blocking benign content. E1 mentioned, *'The EU Commission decided to force the .eu registry to use a list of keywords, and any domain name that contains those keywords has to be sent for extra examination, the list of keywords include words like virus, corona, covid, covid-19, vaccination, vax, anti-vax, there was a whole list, so perfectly innocent websites saying let's say: help covid-19 victims or whatever... completely innocuous, would have been blocked, so it is crazy'*.

E5 noted that legal grounds filtering should be included as a category because it essentially involves listening to court orders. Contrarily, E8 emphasizes that depending on where their business is headquartered, some resolvers may simply choose not to abide by court rulings. Intellectual property filtering was highlighted as contentious by three experts because it can be avoided in any case, just like legal filtering. While E4 recommended using several lists to filter domains to check for overlap and avoid mistakes.

PDNS vs other security measures. Four experts described the main difference between PDNS and other countermeasures as that with DNS is possible to block the source of the problem and once the DNS path is broken DNS abuse will be stopped. For instance, one of these experts claimed that while DNS cannot be bypassed, encrypted network traffic can totally flow through firewalls.

E6 added that with this countermeasure is possible to detect patterns of malicious queries without any indicator of compromise. For instance, a system or user device may be investigated if it increases DNS requests to a domain, even if it's not immediately evident that this is harmful. E4 added that PDNS is a protection mechanism for passive users who want protection.

Two experts, mentioned that PDNS is a solution easy to deploy and that can protect any

device connected to a network without installing anything in each particular device, even if a device does not have any other tool to protect itself, for instance, Internet of Things devices. *'Homes are filled with IoT stuffs, they are WIFI connected, there are heating controllers, in these devices, there is no way you can install an antivirus or to do checks. They have limited hardware. It is important to look at this at the network level and Internet connection, if you install it [PDNS] there, then it works for any possible device that you connect to your network. This is why the filtering is very different than many other applications'*, E2 mentioned.

E8 compared PDNS with the browser, highlighting that it's unusual for DNS to prevent something the browser didn't. However, for devices that do not use the browser, such as the Internet of Things, this might be the only solution available. E8 also mentioned that this is the cheapest solution to deploy.

Privacy. Seven experts talked about how a system like PDNS can affect privacy. According to E1, E2, and E5 the General Data Protection Regulation (GDPR) may apply to these services if they operate under European Union (EU) privacy rules. Other services outside the EU are exempt. E1 said that DNS information might potentially be sold for marketing purposes *'In the EU, you got GDPR, so you got some level of protection if actually, companies are complying, but theoretically it should not be a problem, it should be a non-issue, I don't know if in practice or not, but in theory. Outside the EU, good luck! There are services that are offered in the West, if you are in marketing, there are places where you can buy DNS data'*. Two experts, noted that there is a privacy issue because some public resolvers are even open about sharing DNS information with outside parties. E3 pointed out that while DNS protection is hard to monetize, providers may turn to data collection as a revenue stream. Another expert (E8) stated that even device vendors redirect DNS traffic to them in order to 'guard users' privacy', which really means that they have access to the data.

Transparency. Given the implications that PDNS services have regarding privacy and blocking content, transparency was a topic that experts brought up. Six experts discussed the need for some level of disclosure. *'There must be an understanding of what you are blocking and why you are blocking'* E1 said.

Centralization. Six experts agreed DNS centralization is a problem. Many users relying on one resolver create single points of failure. E4 stated *"I think it [DNS centralization] is creating single points of failure, for me both in professional and personal perspective is a significant reason not to do it ...than the ideological reasons that you will have people arguing on both sides ..."*. E4 also said customers need a variety of options. According to E3 if ISPs provided PDNS to their clients, centralization might be avoided while yet reaping the benefits. E8 stated that there are no security observations that can be made if all DNS

traffic goes to the same party.

Overall, experts emphasized that users must have a thorough understanding of PDNS policies, including what is blocked, how their data is handled, and what they are safeguarded against. Users, ISP customers, enterprises, and experts appear to agree that privacy and transparency must be taken into account in PDNS adoption. Experts suggest that enterprises should think about PDNS performance, the size of the enterprise, the organization's strategy, and network requirements. This implies that some perceived usefulness is important and can vary depending on the type of enterprise. This is consistent with the views of enterprises' participants. Participants in all our previous studies mentioned trust in the provider as critical, and some experts agree.

6.5. DISCUSSION

Our four human subjects studies show how diverse factors can hinder or encourage PDNS adoption. Even though some participants in the ISP interviews said they had extra security measures and did not find PDNS useful, those who used their ISP's PDNS found it useful. This correlates with the survey results, which show that as the perceived usefulness of the service to handle the most important threats for participants increases, so does its adoption. Consistent with [9], secure tools must be useful to be accepted. Also, the trigger – who provides the service – played an important role. Experts emphasized the importance of user awareness. If users are not aware of the conditions of the service and how different it is from other security countermeasures, it is likely that they cannot perceive its utility.

Enterprises' PDNS adoption motivations vary. Enterprises that adopted PDNS did so because of global threat intelligence or an in-depth defense strategy. ISPs offering PDNS to their customers believed security was important and monetize the service after complying with government regulations. Costs and time to handle false positives were elements of simplicity [123] mentioned by participants. To consider adoption other factors such as organizational values, service effectiveness, and how PDNS complements its current infrastructure were mentioned by enterprises' participants. Experts concurred with all these factors adding that knowing where the data of the enterprise is going is important to consider. Some experts also shared the same opinions as users about PDNS providers, including ISPs as potential providers, although highlighting the lack of incentives for this actor to offer PDNS.

6.5.1. INTENTION VS BEHAVIOR

According to [256], there is a chasm between intention and behavior, and we observe discrepancies between the Prolific survey and ISP customer interviews. Perceived vulnerability and data being stolen made PDNS adoption more likely for the survey participants. Yet

only 9% of ISP customers opted in for the ISP PDNS, even though they had suffered a malware infection three months prior to the interview. A possible explanation is that the ISP's message lacked the 'spark' that would have motivated customers to adopt the service. Trigger moments at the correct time can encourage behavior [217], and interventions that encourage progress monitoring may be more successful [256]. However, the ISP did not follow up with a reminder to customers to enable the service.

The Prolific survey found that users who had additional countermeasures in place, namely antivirus, firewall, and ad blockers were more likely to adopt PDNS. However, using these security countermeasures was cited by 27% of ISP customers who did not sign up for the ISP PDNS. We explained how DNS worked, so survey participants may have understood that this was a different countermeasure. In the ISP customer interviews, consumers may not have understood the main value of PDNS because the ISP did not distinguish it from other solutions. Seeing PDNS as a different security countermeasure may have influenced the decision to adopt the service [49].

6.5.2. PROS AND CONS OF PDNS BY DEFAULT

Like every other technology, PDNS offers both benefits and drawbacks, necessitating moral reflection on their use [245]. A PDNS service can detect threats which individual users may never be aware of. This advantage was recognized by some of the survey respondents which stated that ISP as the provider would have the most understanding of current threats, (49 participants), though a few of the ISP interviewees did not immediately recognize this distinction. By enabling PDNS by default users won't have to worry about protecting their devices from dangers that are, for the most part, invisible to them. Our results indicate that defaulting to PDNS may be welcomed by users who perceive it as protecting them, safeguarding their privacy, and being effective in achieving these goals.

However, defaulting users to PDNS might have drawbacks. Privacy issues may arise in jurisdictions without privacy-preserving regulations when inspecting DNS queries. Participants in different studies expressed privacy concerns, and experts highlighted that DNS data may be used by some PDNS providers for commercial purposes. Thus, forcing this countermeasure might not please privacy-conscious users.

Only 9% of the invited ISP customers opted in for PDNS even after a malware infection. Some users preferred to manage their own security or did not consider PDNS effective for their circumstances. Dodier et al. [91] show how ignoring users' priorities and values can result in users circumventing security measures or refusing to implement them. Hence, enforcing PDNS might be counterproductive for these types of users since they might adopt riskier behaviors to trespass DNS blocking or users may adopt self-censoring behaviors [49].

Our findings suggest that enterprises may oppose DNS blocking if they are forced

to employ it by default. Although the organization representatives spoke variously of advantages, many disadvantages were also cited; deploying PDNS would not necessarily negate those concerns.

6.5.3. GOVERNMENT INITIATIVES

Only 8% of the survey participants had a favorable opinion of the government as PDNS provider, while 53% participants preferred their ISP, follow by 34% preferring a commercial company. When asked which provider they would not choose, 65% said the government. These findings suggest PDNS initiatives may be misguided. Our findings suggest that if governments want to stimulate PDNS use, they need to provide users with different alternatives. Users prefer ISPs as providers, so government resources can be directed to involve this actor.

Cost is one of the factors that organizations and survey participants seem to consider for adoption. Government initiatives that are free can be an advantage over commercial providers. Organizations that adopted commercial providers highlighted the importance of global threat intelligence, a capability that local governments might not be able to offer. Service level agreements, efficiency, and effectiveness of the service are among the factors organizations consider for adoption. To match commercial providers, government initiatives may have to compete.

These findings also imply that before proposing user-facing security technologies, user needs and adoption factors should be assessed. Also, determining who users prefer as the intervention's 'facilitator' to adopt a security behavior is crucial. This can determine the success or failure of the behavior.

6.5.4. RECOMMENDATIONS

From our analysis and results, we propose the following recommendations.

Increase PDNS visibility for users. Our findings demonstrate that 71 out of the 121 participants who had never heard of PDNS describe the service as useful. So, users who need the service may not know it exists. Positioning solutions to be easily found by those who can use them is then a challenge. In addition, those providing the services should communicate the benefits and differences of PDNS to end users, distinct from other security measures.

Subsidizing ISPs. Few users and enterprises want their government PDNS. Governments could support ISPs in offering PDNS as our participants mostly saw the ISP as best-placed and prepared to manage such a solution. Subsidizing ISP DNS software and staff is one option. In this way, the PDNS alternative can be available to their customers. A remaining challenge that is pointed out in subsection 6.5.5 is to investigate ISPs' incentives to actually offer PDNS.

Blocklist sharing. Sharing blocklists with PDNS resolvers as other existing abuse data is shared (e.g. as Shadowserver does [255]) might be an alternative for governments. Enterprises can subscribe to receive these blocklists and implement them in the way they see fit. Two enterprises deployed their own in-house PDNS and supported this. This can provide another approach to deploying PDNS, while complementing options for paid-for commercial threat intelligence.

6.5.5. LIMITATIONS AND FUTURE WORK

We calculated PDNS penetration using the ‘first use’ resolver from APNIC data. Thus, our results are a lower-bound estimate of PDNS penetration since we do not include ‘all resolvers’ that may view users’ DNS requests.

Instead of a random sample of the general population, we recruited survey respondents using Prolific. Tang et al [272] suggest that Prolific data is representative of user views and experiences. We recruited participants from different regions, so views are not localized.

We interviewed participants who work in 12 enterprises in government, banking, university, cable industry, and ISPs. We found recurring themes that drove or hindered PDNS adoption. Why an organization adopts PDNS might vary, but more organizations’ viewpoints do not invalidate our findings.

This work focuses on public DNS resolvers advertising themselves as PDNS and particularly protecting against botnets, malware, phishing, and spam; other categories like adult content were not considered.

Further research may determine whether ISPs are using PDNS without offering it as a service. Which incentives ISPs have to offer PDNS is also worth exploring as well as privacy trade-offs users might be willing to make. Gender was a control variable in our survey instrument, so more research may need to confirm gender differences in PDNS adoption. We did not find differences in PDNS adoption across regions; however, certain regions in our sample fell below 30 observations, so further research may corroborate these findings.

6.6. CONCLUSION

Using APNIC dataset, we found that commercial PDNS resolvers are marginally used in different regions with Africa having the highest adoption rate at 2%. Nonetheless, some countries, with Israel on the lead, have a high rate of adoption. Four human studies identified PDNS adoption factors for users and organizations. Perceived vulnerability, perceived usefulness, and cost played an important role in users’ adoption. Enterprises used PDNS for global threat intelligence, layered security, believing security was important, and compliance with government regulations. Also, our findings demonstrate the need of considering user

preferences for intervention facilitators when recommending user-facing security solutions like PDNS.

7

CONCLUSION

This dissertation studied the role of users, manufacturers, and intermediaries in IoT security. We have presented five peer-reviewed studies (see Chapters 2 - 6), that aimed to respond the following main research question:

How can users mitigate infected IoT devices? And what role can manufacturers and intermediaries play in supporting them?

In this final chapter, we look back on the work presented in the previous chapters (section 7.1). We also discuss how the five peer-review studies answer our main research question (section 7.2). Next, we consider the practical governance and policy implications of our findings (section 7.3). Finally, we discuss potential research directions following up the findings in this dissertation (Section section 7.4).

7.1. SUMMARY OF THE FINDINGS

7.1.1. CHAPTER 2: USER COMPLIANCE AFTER IoT MALWARE NOTIFICATIONS

As we saw in chapter 2, there are usability issues with informing consumers about an infected IoT device on their network. Users could not be informed of which device(s) had been compromised with certainty. At least for this European ISP located in The Netherlands that

is bound by regulations like the General Data Protection Regulation (GDPR) and concepts like net neutrality. Thus, the ISP's recommendations to customers were generic in nature. We demonstrated that compliance with the recommended steps, even if they were generic, increased the probability of cleanup of infected IoT device(s) by 32%. After being notified, users were eager to follow the instructions. Also, we learned that the notification acted as an attention switch, causing users to take action. These findings highlighted the importance of the intermediary intervention. In addition, we learned that users' motivations such as malfunctioning of the device or just needing the device had a negative impact on compliance compared to users who wanted their Internet back since they were quarantined (without an Internet connection).

7.1.2. CHAPTER 3: REAL-WORLD INTERVENTIONS IN SMART HOME SECURITY

In chapter 3, we used think-aloud observations to learn the process that took place in users' homes to capture users' efforts to comply with an intermediary notification regarding an infected IoT device(s) in their network. Our research showed that users had difficulty following the suggested remediation steps. Given the general nature of the notification, the initial challenge was identifying the culprit device(s). We found evidence that some manufacturers do not provide appropriate support for users to change passwords (for instance, some users described that manuals or manufacturers' websites do not provide useful information to carry out this task), users anticipated some feedback or confirmation on the success of performing the steps (which they could not obtain from the device(s) itself or the ISP as the sender of the notification), and users resorted to familiar behavior such as disconnecting the device(s) to reset it. Users were motivated to take the necessary actions after receiving the notification even though they took time and effort to 'clean up the device(s)'.

7.1.3. CHAPTER 4: REMEDIATING PERSISTENT IoT MALWARE

In chapter 4, we explored users' efforts in remediating a persistent IoT malware, QSnatch. We partnered again with an intermediary to notify users and compare the remediation time of this malware family against Windows and non-persistent IoT malware. QSnatch had a survival probability of 30% after 180 days of a device being compromised; while most, if not all, other malware infections were removed. Hence, our results confirmed that persistent IoT malware takes longer to remediate than Windows and non-persistent IoT malware. We encountered an ideal scenario in which the vast majority of users possessed technical competency and did not find the remediation process difficult, displayed motivation to complete the steps, but required time to do so and multiple notifications before taking any

action. Once again, this research provided evidence that notifications are crucial. It was also emphasized that users cannot always be provided with detailed information about the compromised device type and tailor-made advice for resolving the issue.

7.1.4. CHAPTER 5: IOT MANUFACTURERS' ROLE IN DEVICE INFECTIONS

In chapter 5, we identified manufacturers of devices that get most often compromised with Mirai-like infections. We found that only 9 vendors shared almost 50% of the infections. These findings implied that encouraging a more security-minded posture from this particular subset of manufacturers could improve IoT security as a whole. Also, we found evidence that 53% of these manufacturers offer firmware or software updates on their websites, 43% of these manufacturers offer password-changing procedures information on their security advisory websites, and 26% of the manufacturers provide recommendations for protecting the devices from attacks. These findings suggested that providing security updates and advice does not prevent or remediate infected IoT devices. A possible reason for this finding is that the advice provided by manufacturers is not actionable or hard to implement, and it is fragmented.

7.1.5. CHAPTER 6: UNDERSTANDING PROTECTIVE DNS ADOPTION FACTORS

In chapter 6, we explored 'Protective DNS' (PDNS) adoption. A service that leverages DNS resolutions to prevent malicious activity in users' networks, including IoT. According to our research, protective DNS resolvers offered by commercial companies have a low adoption rate, yet in some countries, more than 20% of DNS queries are routed through them. We undertook the task to understand factors that would drive users' adoption of PDNS. Our Prolific survey revealed that users would adopt PDNS if offered by their ISPs over their governments, and their second-best choice would be a commercial company. Through interviews with real-world ISP customers who were offered a free PDNS service after a malware infection, we found that only 9% of the customers opted in for the service, and some of the reasons listed for not adopting the service were not needing the service (perhaps because they did not understand the difference between the service and security countermeasures like antivirus), being unable to turn on the service, fear of being charged later and the service blocking something that they would like to access. These findings highlighted that users might benefit from seeing this as a distinct option from any other security countermeasure. Among enterprises, we found that 50% of the participants were already using PDNS. The primary drivers of adoption were access to global threat intelligence, the implementation of a layered security approach, and the fulfillment of regulatory requirements. Experts highlighted the

limitations of PDNS and the broad challenges of implementing this countermeasure such as transparency and centralization.

7.2. DISCUSSION

In the introduction, we noted that information asymmetry is present in the IoT market. Furthermore, misaligned incentives of the actors involved are major problems for IoT security. Stakeholders have different power and interest in solving infected IoT devices. The inadequate security of IoT devices causes negative externalities for third parties. Third parties, who are neither the buyer nor the seller of these devices, must protect themselves against Distributed Denial of Service (DDoS) attacks. No easy solution can be provided to such a ‘wicked’ problem (infected IoT devices), so it is possible that a greater outcome could be achieved by a concerted effort on the part of the various stakeholders involved.

We turn to answer our main research question **How can users mitigate infected IoT devices? And what role can manufacturers and intermediaries play in supporting them?** by reflecting on the role of each stakeholder namely users, IoT manufacturers, and intermediaries in light of our findings.

7.2.1. USERS

Chapters 1-3 demonstrate users’ motivation to remediate infected IoT devices once they were notified. Especially in chapter 3, we observed the effort users make to comply with the recommended steps.

A problem that crops up in both chapter 2 and chapter 3 is that it is not always possible to provide users with tailor-made advice or inform them exactly which device is the one infected. Advice, according to the literature, should be actionable [235]. If a behavior is to be put into practice, it must be simple to do so and incur no additional costs [123]. Also, the literature suggests that security is not the primary task of users and technology should be designed in a way that does not burden users [248]. The time and effort required to implement the behavior of ‘remediating an infected IoT device(s)’, when users are given generic guidance, are high.

Even though there is no legal requirement for ISPs to notify users about infected IoT devices, some ISPs have undertaken this task. This is a voluntary action. However, ISPs cannot provide feedback or confirmation of the success of the steps that they are recommending. Besides not having a feedback loop of success in their efforts, users face a lack of support from websites or manuals of manufacturers as was shown in chapter 3. In the absence of any user-informing tools, ISPs’ best efforts are a starting point to urge users to act.

By helping to clean up infected devices, users can play a significant role in reducing

the potential size of IoT botnets. Nonetheless, users cannot get rid of compromised IoT devices without being aware of them and without actionable advice. These two aspects are key to empowering them to solve the security issue. Some manufacturers are addressing usability concerns by providing malware removal tools (like QNAP in chapter 4), which is encouraging. Unfortunately, the advice users encountered is convoluted, so there is still more to be done to eradicate the widespread usability problems that users face dealing with infections.

When users are informed by a third party (their ISP) that an infected IoT device is in their network, the results of chapters 2-4 suggest that they are doing what they can to remediate the infected IoT devices using the advice that is given to them, the tools and familiar behaviors that are at their disposal. As soon as they are made aware of the issue, users are driven to take action and clean up the infected IoT device(s).

7.2.2. MANUFACTURERS

In chapter 5, we observe that the abuse of IoT devices is concentrated in the hands of a few companies. Approximately half of all Mirai-like infected IoT devices are associated with just nine IoT manufacturers. Mirai is malicious software that takes advantage of weak or default login credentials. Hence, these results have clear implications for how IoT manufacturers can assist in reducing IoT botnets; IoT manufacturers can help mitigate IoT botnets by removing from the setup process of their devices easy-to-guess passwords or default credentials and replacing them with unique strong passwords. Even though in chapter 5, we found evidence of IoT manufacturers offering password-changing procedure advice and software and firmware updates, this does not prevent infected IoT devices. Once the device is in use, users might never login again to configure a different credential or they could forget about it. Thus, in order to facilitate this process for users, IoT manufacturers could ensure that this is already the default option or that users cannot use their devices unless the password has been changed. Of course, this has also some implications for usability since users might get frustrated if they cannot use their IoT device if the process of changing the password is too complicated. Thus, IoT manufacturers should consider simple password-changing processes if they opt for this option.

Several best practices, such as the European Union Agency for Cybersecurity (ENISA) standards for IoT software development and the United Kingdom's Code of Practice for consumer IoT Security, are in place to encourage manufacturers to secure IoT devices. But most of these measures are adopted voluntarily or not at all by manufacturers. Some countries such as the United Kingdom have opted already for a more strict approach banning the use of default credentials via regulatory approach and other countries are moving towards

that direction [182]. However, the implementation of these laws and regulations can take few years. Thus, if this handful of manufacturers changes this basic security practice, the outcome can have a clear impact on IoT security.

Also, manufacturers can provide users with less fragmented and actionable advice via their security advisory. As we observed in chapter 3, users encounter a lack of support from manuals and websites. Even if future laws or standards get rid of default credentials, legacy products will remain in the market. Some of the websites of manufacturers focus on presenting the features of the products, comparing costs, thus this highlights room for improvement in how manufacturers present security advice to users to change passwords or other security procedures.

7.2.3. INTERMEDIARIES

Intermediaries support the fundamental platforms and infrastructure of the Internet and facilitate communications and transactions between third parties and services [219]. Cetin et al. [55] showed that 80% of the infected IoT devices are in broadband networks, so they are in a privileged position to intervene in this problem. In chapter 2, chapter 3, and chapter 4, we see the important role of ISPs notifying users about an infected IoT device in their network. Without the external intervention of this actor, users might not be even aware of the security issue. Particularly in chapter 4, we observed that participants needed multiple notifications in order to act.

Third parties are frequently used by both producers and consumers to help address information asymmetry; one way they can help is by providing unbiased information about a good [285], in this case, IoT devices. For instance, users can subscribe to consumer reports that actually test products and publish their results [107]. Along these lines, ISPs can help reduce infected IoT devices by offering a notification subscription service. ISPs can gain users' consent to point out the culprit device using network scans and provide tailor-made advice to users. In this way, intermediaries do not only provide timely information to users to act towards the security issue, but also users learn about the quality of the goods they bought. In future purchases, users might consider their experience with certain brands, which in turn can provide manufacturers some incentives to improve their security and support.

It is important to consider that in contrast to personal computers, IoT devices lack computational capacity and battery power. As a result, standard protection strategies, such as signature-based anti-virus software, cannot be used to protect the heterogeneity of IoT devices present in the market [85, 97, 294] or even if manufacturers want to notify users this might not be possible in all cases. Thus, incentivizing ISPs notifications can contribute to mitigating IoT botnets.

In chapter 6, we explored the possibility of a protective DNS service to prevent malicious activity in users' networks. Users distrust the government for this task, and this prevention mechanism can succeed if ISPs offer it to them. Also, ISPs are one of the main DNS providers for users, so they could offer a prevention service that users can subscribe to and where ISPs can report to the user the malicious activities that they spot and from which devices and brands, this, in turn, can aid users to make future decisions in their IoT purchases, while also preventing infections.

After reflecting on the roles of each actor, **How can users mitigate infected IoT devices? And what role can manufacturers and intermediaries play in supporting them?** Going back to the 'Power-interest' grid [10] presented in section 1.6, users can be moved to the 'Players' quadrant by providing them with notifications via intermediaries (which reduces information asymmetry about infections as well as facilitates the process of threat detection in users network) since our research shows that they have the motivation to act, but they lack the opportunity to identify the threats and they struggle with generic advice. Improving notifications with tailor-made advice can make remediation easier, but assuming tailor-made advice is not feasible, even general recommendations with a notification are preferable than users being in the dark about the security of their IoT devices. To support users, manufacturers should at least remove easily guessable credentials from the setup process of their devices. It is a handful of manufacturers that could implement this change and they will have a large impact on IoT security. The evidence suggests that users are eager to assist in the mitigation of malicious activity in their networks by adopting a preventative service offered by their Internet service providers (ISPs) that makes use of DNS. Thus, intermediaries also can support users by offering such services. However, this requires that ISPs are willing to move to the 'Player' quadrant.

7.3. IMPLICATIONS FOR GOVERNANCE AND POLICY MAKING

The term 'governance' refers to the collective efforts taken by many groups to address social problems, including but not limited to governmental agencies, public bodies, the private sector, and civil society [184]. According to Meuleman [184] there are three ideal forms of governance: the hierarchical form, the market form, and the network form. Under a hierarchical structure, the government is seen as an integral part of the problem-solving process, and public responsibility is emphasized. Market governance is a bottom-up, incentive-based approach that rejects the rigidity of hierarchical structures in favor of voluntary transactions. While network governance is based on the interdependent, trust-based relationships among the actors involved.

Infected IoT devices show the limits of the market governance style. When it comes

to DDoS attacks, neither manufacturers nor users suffer any consequences of the poor security of IoT devices. Third parties incur the cost of protecting themselves against DDoS attacks, creating a negative externality. IoT market also suffers from Information asymmetry. Users have less information about the security state of their IoT devices than manufacturers. Negative externalities and information asymmetry are known as market failures, and they justify government intervention [285]. Thus, this raised questions about the role of governments in defining laws, rules, and regulations for IoT security in their countries [270], or a hierarchical governance style. However, since the government is not an all-powerful actor, and this is a network of stakeholders, attempting to implement a policy or strategy without the help of the actors involved may be doomed to failure [80]. Given the complexity of the problem and the need to balance competing interests of stakeholders in remediating infected IoT devices, a hybrid approach of hierarchical governance, network governance, and market governance may be the best way to balance the IoT market. We will discuss each approach individually while keeping in mind that a combination of them may be ideal.

7.3.1. HIERARCHICAL GOVERNANCE

Some governments have attempted to solve IoT security issues by releasing guidelines for secure IoT devices, [115, 150], good practices for secure IoT software development [104], and encouraging transparency of the whole supply chain of IoT devices [5]. Others such as Singapore, Finland and Germany voluntarily have adopted labels [8]. Some regions such as the European Union are imposing strong rules on certifications [105, 291]. All these efforts in the long run can help to reduce the influx of insecure IoT devices in the market.

However, there is a variety of rules and regulations that apply to different countries. Hence, there are jurisdictional issues that are difficult to address. For instance, users might buy devices online (there are many platforms such as Amazon where products can be bought online today). Governments might not want to restrict certain products since this can lead to black markets emerging [285] and insecure devices will be present in the market anyway. Moreover, commercial trading relationships among countries might be damaged if products are restricted.

Our results point to a concentration of manufacturers that get most often compromised due to the use of default credentials. Governments can direct resources (which in some cases might be limited) to engage these manufacturers in more secure practices – particularly, removing from the setup process of IoT devices the use of default and easy-to-guess credentials.

Some certification entities within countries' jurisdiction that have no interconnected

interest with manufacturers can regulate that minimum IoT security requirements are ensured. For instance, in the European Union, there are National Cybersecurity Certification Authorities (NCCAs) [105], and they call for banning the use of default passwords in IoT devices among other basic security measures [188]. Enforcement and accountability by certification authorities – particularly of no default credential use – could reduce the high amount of IoT device infections. Even if not all governments have the capacity to create these certification bodies, regional initiatives could lead to potential benefits for everyone. If manufacturers have to remove from their practices the use of default and easy-to-guess credentials as the bare minimum to ship products to Europe, it might result easier for IoT manufacturers to apply these requirements to all the devices they produce.

Governments can ask also for strong warranties [285], which can protect users from buying a device that within a few months might not have support to solve security issues and that generates security costs to the user. Returning policies and retailers' involvement can be part of strong warranties. Users can approach retailers easily, contrary to manufacturers. In our results, we observed that most often compromised manufacturers were located in Asia. Thus, initiatives such as the one in The Netherlands, where retailers are responsible for updates [189, 211], are interesting since manufacturers of IoT devices are located all over the world. Being able to return devices can incentivize manufacturers to produce secure products with the necessary support for users.

7.3.2. NETWORK GOVERNANCE

Since many IoT products make it to market without implementing even the most basic security measures, such as not using easy-to-guess passwords, laws and regulations to achieve the bare minimum in security are necessary. However, a network governance model that includes stakeholders and ensures the government works with them as partners rather than imposing rules and regulations, can yield better results in some circumstances [184]. Considering the diverging and varying interests of the stakeholders, network governance can be adaptive and facilitate that stakeholders organize themselves.

ISPs, as intermediaries, have access to millions of users and they can see their network traffic. The government could consider subsidizing ISPs' notification activities. CyberGreen [75] proposes a cyber-public health approach to deal with cybersecurity issues. Notifications can be seen as a way to create awareness to deal and prevent future infections, as it was done with some deadly diseases in the past. Another option is that ISPs could gain users' consent to scan users' networks to offer notification services and provide tailor-made advice. This can generate stream revenues for ISPs, and this might incentivize this stakeholder to endeavor this task. Also, some synergy could be encouraged by governments between IoT

manufacturers and ISPs, so IoT manufacturers could partner with ISPs in order to notify users.

From chapter 6, we learned that users are willing to adopt a protective DNS service to prevent malicious traffic, if offered by their ISPs. Governments can consider subsidizing ISPs to implement such prevention methods. Another option is that since the interest of Internet users in ISP-based security services has been measured [246], and similar in chapter 6, 41% of the surveyed users are willing to pay for security. A combination of economic incentives for ISPs offering these services and governments encouraging these actors to part take in the solution can also provide a starting point for this alternative prevention mechanism.

Governments can also pursue awareness campaigns for users. Even if users adopt preventive services such as protective DNS, these services might not be always perfect. Thus, users need to be able to understand that IoT devices can be compromised and the costs for society, and how they could be able to adopt secure behaviors.

7.3.3. MARKET GOVERNANCE

Despite the fact that we have referred to IoT security as a market failure, market dynamics, in particular healthy competition among IoT manufacturers, may result in IoT devices of improved quality. Innovation is fueled by competition, and in this scenario, manufacturers can set standards for the minimum expectations that customers can have when purchasing one of their devices. As we can see in chapter 4, for instance, QNAP provides the ability to run a malware removal tool, which may encourage other manufacturers to compete by offering similar solutions if the device capacity allows it.

Transparency and data-driven reports as well as bench-marking IoT manufacturers and products can also improve IoT security. Similar to the work presented in chapter 5, information availability may have a reputation impact on IoT manufacturers, so scholarly efforts are always valued. Initiatives of countries such as The Netherlands, where the consumer association test and make public to users the flaws of IoT devices, can incentivize manufacturers' action [107]. If there is more than one interaction between users and these IoT manufacturers' brands, users might choose not to buy them in the future. Transparency can make possible for the demand (users) of these products to decide by themselves which products they want to buy [302].

To reduce information asymmetry manufacturers could advertise that they do not use easy-to-guess passwords or default credentials in their products, so users can make more informed decisions when purchasing the devices. This is different from the proposed labels to assess the security of a device, but more like a self-marketing strategy that can be included in the box of IoT products, which is more simple to attain. This would not require any

assessment of any body, but IoT manufacturers could signal users of potential basic security.

7.4. FUTURE WORK DIRECTIONS

Each study's shortcomings are addressed in detail in their respective chapters. Three avenues exist for further research.

First, studying usability issues regarding advice to handle infected IoT devices. In chapter 2 and chapter 3, the advice provided to users was generic, while in chapter 4 was tailor-made. We did not compare in any of our studies how difficult or easy it is for users to perform the remediation when the advice is tailor-made versus when users have to deal with uncertainty, such as identifying the devices themselves. Also, an exploration of how to make advisory manufacturers' pages more user-friendly and their advice more actionable can be researched. In addition, there are different sources of IoT security advice that can be studied, for instance, government advice or retailers' advice.

A second area of future research is the incentives of ISPs on offering notifications and protective DNS services. The studies of this dissertation were carried out with ISPs that already notify users and offered PDNS. Hence, understanding how to involve ISPs who are not implementing notifications is crucial as well as learning if ISPs are willing to offer a PDNS service and what factors they would consider for providing it to its customers.

Finally, different types of ways to deliver signals for IoT security and diminish the information asymmetry for users when buying a device can be studied. There is some work studying security labels, but there are other ways to signal security to users. For instance, signaling security via retailers' websites, store stands, or e-commerce websites selling IoT products.

BIBLIOGRAPHY

- [1] [n. d.]. GitHub - sensepost/gowitness: gowitness - a golang, web screenshot utility using Chrome Headless. <https://github.com/sensepost/gowitness>
- [2] [n. d.]. GitHub - zmap/zgrab2: Fast Go Application Scanner. <https://github.com/zmap/zgrab2>
- [3] [n. d.]. I scanned the whole country of Austria and this is what I've found. <https://blog.haschek.at/2019/i-scanned-austria.html>
- [4] [n. d.]. Quad9 | A public and free DNS service for a better security and privacy. <https://www.quad9.net/>
- [5] [n. d.]. Software Bill of Materials | National Telecommunications and Information Administration. <https://www.ntia.gov/SBOM>
- [6] [n. d.]. TeleGeography | Home. <https://www2.telegeography.com/>
- [7] 2021. The Shadowserver Foundation. <https://www.shadowserver.org/>
- [8] 2023. Are Consumers Willing to Pay for Security and Privacy of IoT Devices?. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA. <https://www.usenix.org/conference/usenixsecurity23/presentation/emami-naeini>
- [9] Ruba Abu-Salma, M Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. 2017. Obstacles to the adoption of secure communication tools. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 137–153.
- [10] Fran Ackermann and Colin Eden. 2011. Strategic Management of Stakeholders: Theory and Practice. *Long Range Planning* 44, 3 (jun 2011), 179–196. <https://doi.org/10.1016/j.lrp.2010.08.001>
- [11] George A Akerlof. 1978. The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*. Elsevier, 235–251.

- [12] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*. USENIX, Washington, D.C., 257–272. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe>
- [13] Fadele Ayotunde Alaba, Mazliza Othman, Ibrahim Abaker Targio Hashem, and Faiz Alotaibi. 2017. Internet of Things security: A survey. *Journal of Network and Computer Applications* 88 (2017), 10–28. <https://doi.org/10.1016/j.jnca.2017.04.002>
- [14] D Allen. 1999. Transaction Costs. *Encyclopedia of Law and Economics* (1999).
- [15] Kevin Allix, Quentin Jérôme, Tegawende F Bissyandé, Jacques Klein, Radu State, and Yves Le Traon. 2014. A Forensic Analysis of Android Malware—How is Malware Written and How it Could Be Detected?. In *2014 IEEE 38th Annual Computer Software and Applications Conference*. IEEE, 384–393.
- [16] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monroe. 2019. SoK: Security evaluation of home-based IoT deployments. In *2019 IEEE symposium on Security and Privacy (S&P)*. IEEE, 1362–1380.
- [17] Omar Alrawi, Charles Lever, Kevin Valakuzhy, Ryan Court, Kevin Snow, Fabian Monroe, and Manos Antonakakis. 2021. The Circle Of Life: A Large-Scale Study of The IoT Malware Lifecycle. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3505–3522. <https://www.usenix.org/conference/usenixsecurity21/presentation/alrawi-circle>
- [18] Rosemarie Anderson. 2007. Thematic content analysis (TCA). *Descriptive presentation of qualitative data* (2007), 1–4.
- [19] Ross Anderson and Tyler Moore. 2006. The Economics of Information Security. *Science* 314, 5799 (oct 2006), 610–613. <https://doi.org/10.1126/SCIENCE.1130992>
- [20] Sandhya Aneja, Nagender Aneja, and Md Shohidul Islam. 2018. IoT Device Fingerprint using Deep Learning. In *2018 IEEE International Conference on Internet of Things and Intelligence System (IOTAIS)*. IEEE, 174–179. <https://doi.org/10.1109/IOTAIS.2018.8600824>

- [21] Eirini Anthi, Shazaib Ahmad, Omer Rana, George Theodorakopoulos, and Pete Burnap. 2018. EclipseIoT: A secure and adaptive hub for the Internet of Things. *Computers & Security* 78 (2018), 477–490.
- [22] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. 2017. Understanding the Mirai Botnet. *Proceedings of the 26th USENIX Security Symposium* (2017), 1093–1110. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis>
- [23] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, and others. 2017. Understanding the mirai botnet. In *USENIX Security Symposium*. 1092–1110. <https://doi.org/10.1016/j.religion.2008.12.001>
- [24] Arman Noroozian, Elsa Rodríguez, Elmer Lastdrager, Takahiro Kasama, Michel van Eeten, and Carlos Gañán. 2021. Can ISPs Help Mitigate IoT Malware? A Longitudinal Study of Broadband ISP Security Efforts (to appear). *IEEE Euro Security & Privacy* (2021).
- [25] Todd W Arnold. 2010. Uninformative parameters and model selection using Akaike’s Information Criterion. *The Journal of Wildlife Management* 74, 6 (2010), 1175–1178.
- [26] Hadi Asghari, Michael Ciere, and Michel J G Van Eeten. 2015. Post-Mortem of a Zombie: Conficker Cleanup After Six Years. *USENIX Security* (2015). <https://doi.org/10.1175/jhm-d-12-0146.1>
- [27] Hadi Asghari and Arman Noroozian. 2022. GitHub - hadiasghari/pyasn: Python IP address to Autonomous System Number lookup module. (Supports fast local lookups, and historical lookups using archived BGP dumps.). <https://github.com/hadiasghari/pyasn>
- [28] Hadi Asghari, Michel J.G. van Eeten, and Johannes M. Bauer. 2015. Economics of Fighting Botnets: Lessons from a Decade of Mitigation. *IEEE Security & Privacy* 13, 5 (sep 2015), 16–23. <https://doi.org/10.1109/MSP.2015.110>
- [29] Hadi Asghari, Michel JG van Eeten, and Johannes M Bauer. 2015. Economics of fighting botnets: Lessons from a decade of mitigation. *IEEE Security & Privacy* 13, 5 (2015), 16–23.

- [30] Kevin Ashton et al. 2009. That ‘internet of things’ thing. *RFID journal* 22, 7 (2009), 97–114.
- [31] Terrence August, Duy Dao, and Kihoon Kim. 2019. Market Segmentation and Software Security: Pricing Patching Rights. *Management Science* 65, 10 (oct 2019), 4575–4597. <https://doi.org/10.1287/mnsc.2018.3153>
- [32] Avtech. 2014. AVTECH - Leader in Push Video HDCCTV, IP Camera, CCTV camera, DVR, IVS Network camera, EagleEyes mobile surveillance, NVR, NAS and CMS total solution. <https://www.avtech.com.tw/EOL.aspx>
- [33] Catherine Barrett. 2019. Are the EU GDPR and the California CCPA becoming the de facto global standards for data privacy and protection? *Scitech Lawyer* 15, 3 (2019), 24–29.
- [34] Christopher Bellman and Paul C van Oorschot. 2020. Best Practices for IoT Security: What Does That Even Mean? *arXiv preprint arXiv:2004.12179* (2020).
- [35] Bitdefender. 2022. Bitdefender BOX: Frequently Asked Questions | FAQ. <https://www.bitdefender.com/consumer/support/answer/13906/>
- [36] J Martin Bland and Douglas G Altman. 2004. The logrank test. *Bmj* 328, 7447 (2004), 1073.
- [37] JM Blythe and SD Johnson. 2018. The Consumer Security Index for IoT: A protocol for developing an index to improve consumer decision making and to incentivize greater security provision in IoT devices. *IET* (2018).
- [38] John M Blythe, Shane D Johnson, and Matthew Manning. 2020. What is security worth to consumers? Investigating willingness to pay for secure Internet of Things devices. *Crime Science* 9, 1 (2020), 1–9.
- [39] John M Blythe, Nissy Sombatruang, and Shane D Johnson. 2019. What security features and crime prevention advice is communicated in consumer IoT device manuals and support pages? *Journal of Cybersecurity* 5, 1 (2019), tyz005.
- [40] Leon Böck, Nikolaos Alexopoulos, Emine Saracoglu, Max Mühlhäuser, and Emmanouil Vasilomanolakis. 2019. Assessing the Threat of Blockchain-based Botnets. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–11.

- [41] Brennen Bouwmeester, Elsa Rodríguez, Carlos Gañán, Michel van Eeten, and Simon Parkin. 2021. “The Thing Doesn’t Have a Name”: Learning from Emergent Real-World Interventions in Smart Home Security. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*. 493–512.
- [42] Irina Brass, Leonie Tanczer, Madeline Carr, Miles Elsdén, and Jason Blackstock. 2018. Standardising a moving target: The development and evolution of IoT security standards. *IET* (2018).
- [43] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [44] Virginia Braun and Victoria Clarke. 2020. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* (2020), 1–25.
- [45] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.
- [46] C. Bravo-Lillo, L. F. Cranor, J. Downs, and S. Komanduri. 2011. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security Privacy* 9, 2 (March 2011), 18–26. <https://doi.org/10.1109/MSP.2010.198>
- [47] Calvin Brierley, Jamie Pont, Budi Arief, David J Barnes, and Julio C Hernandez-Castro. 2020. Persistence in Linux-Based IoT Malware. *NordSec 2020* (2020).
- [48] Pamela Briggs, Debbie Jeske, and Lynne Coventry. 2017. Behavior change interventions for cybersecurity. In *Behavior change research and theory*. Elsevier, 115–136.
- [49] P. Briggs, D. Jeske, and L. Coventry. 2017. Behavior Change Interventions for Cybersecurity. In *Behavior Change Research and Theory: Psychological and Technological Perspectives*. Elsevier Inc., 115–136. <https://doi.org/10.1016/B978-0-12-802690-8.00004-9>
- [50] Cam Davidson-Pilon. 2022. Estimating univariate models — lifelines 0.26.4 documentation. <https://lifelines.readthedocs.io/en/latest/Survivalanalysiswithlifelines.html>
- [51] Fran Casino, Nikolaos Lykousas, Ivan Homoliak, Constantinos Patsakis, and Julio Hernandez-Castro. 2021. Intercepting Hail Hydra: Real-time detection of Algorithmi-

- cally Generated Domains. *Journal of Network and Computer Applications* 190 (2021), 103135. <https://doi.org/10.1016/j.jnca.2021.103135> arXiv:2008.02507
- [52] Cencys. 2020. Home. <https://censys.io/>
- [53] Centre for the Promotion of Imports Ministry of Foreign Affairs Netherlands. 2022. The European market potential for (Industrial) Internet of Things | CBI. <https://www.cbi.eu/market-information/outsourcing-itobpo/industrial-internet-things/market-potential>
- [54] Orçun Çetin, Lisette Altena, Carlos Gañán, and Michel Van Eeten. 2018. Let Me Out! Evaluating the Effectiveness of Quarantining Compromised Users in Walled Gardens. *Fourteenth Symposium on Usable Privacy and Security* (2018).
- [55] Orçun Cetin, Carlos Gañán, Lisette Altena, Takahiro Kasama, Daisuke Inoue, Kazuki Tamiya, Ying Tie, Katsunari Yoshioka, and Michel Van Eeten. 2019. Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. In *NDSS*. <https://doi.org/10.14722/ndss.2019.23438>
- [56] Orçun Çetin, Carlos Gañán, Lisette Altena, Samaneh Tajalizadehkhoob, and Michel van Eeten. 2019. Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Networks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 326–339.
- [57] Orçun Çetin, Carlos Ganán, Lisette Altena, Samaneh Tajalizadehkhoob, and Michel Van Eeten. 2019. Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Networks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 326–339.
- [58] Orçun Çetin, Carlos Ganan, Maciej Korczynski, and Michel van Eeten. 2017. Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. *Workshop on the Economics of Information Security (WEIS)* (2017).
- [59] Orçun Çetin, Mohammad Hanif Jhaveri, Carlos Gañán, Michel van Eeten, and Tyler Moore. 2016. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity* (2016). <https://doi.org/10.1093/cybsec/tvw005>
- [60] George Chalhouh and Ivan Flechais. 2020. “Alexa, are you spying on me?”: Exploring the Effect of User Experience on the Security and Privacy of Smart Speaker Users. In *International Conference on Human-Computer Interaction*. Springer, 305–325.

- [61] BR Chandavarkar. 2020. Hardcoded Credentials and Insecure Data Transfer in IoT: National and International Status. In *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*. IEEE, 1–7.
- [62] Hyunsang Choi and Heejo Lee. 2012. Identifying botnets by capturing group activities in DNS traffic. *Computer Networks* 56, 1 (jan 2012), 20–33. <https://doi.org/10.1016/J.COMNET.2011.07.018>
- [63] Yi Ting Chua, Simon Parkin, Matthew Edwards, Daniela Oliveira, Stefan Schiffner, Gareth Tyson, and Alice Hutchings. 2019. Identifying unintended harms of cybersecurity countermeasures. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–15.
- [64] C. Cimpanu. 2020. CISA says 62,000 QNAP NAS devices have been infected with the QSnatch malware. <https://www.zdnet.com/article/cisa-says-62000-qnap-nas-devices-have-been-infected-with-the-qsnatch-malware/>
- [65] CIRA Canadian Shield. [n. d.]. CIRA Canadian Shield | Free public DNS for Canadians | CIRA. <https://www.cira.ca/cybersecurity-services/canadian-shield>
- [66] Cisco. 2022. Home Internet Security | OpenDNS. <https://www.opendns.com/home-internet-security/>
- [67] Cloudflare. 2020. Introducing 1.1.1.1 for Families. <https://blog.cloudflare.com/introducing-1-1-1-1-for-families/>
- [68] European commission. 2022. Funding & tenders. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/cef-dig-2021-cloud-dns-works>
- [69] Vincent C. Conzola and Michael S. Wogalter. 2001. A Communication–Human Information Processing (C–HIP) approach to warning effectiveness in the workplace. *Journal of Risk Research* 4, 4 (oct 2001), 309–322. <https://doi.org/10.1080/13669870110062712>
- [70] Andrei Costin and Jonas Zaddach. 2018. Iot malware: Comprehensive survey, analysis framework and case studies. *BlackHat USA* 1, 1 (2018), 1–9.
- [71] Andrei Costin, Jonas Zaddach, Aurélien Francillon, and Davide Balzarotti. 2014. A large-scale analysis of the security of embedded firmwares. In *23rd USENIX Security Symposium (USENIX Security '14)*. 95–110.

- [72] E. Cozzi, M. Graziano, Y. Fratantonio, and D. Balzarotti. 2018. Understanding Linux Malware. In *2018 IEEE Symposium on Security and Privacy (SP)*. 161–175. <https://doi.org/10.1109/SP.2018.00054>
- [73] Francisco Cribari-Neto and Achim Zeileis. 2010. Beta regression in R. *Journal of Statistical Software* (2010). <https://doi.org/10.18637/jss.v034.i02>
- [74] Cyber Security Agency of Singapore. 2022. Cybersecurity Labelling Scheme. <https://www.csa.gov.sg/Programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme/about-cls>
- [75] CyberGreen. 2022. Cyber Public Health - CyberGreen. <https://cybergreen.net/cyber-public-health/>
- [76] Cybersecurity and Infrastructure Security Agency. 2020. Potential Legacy Risk from Malware Targeting QNAP NAS Devices | CISA. <https://us-cert.cisa.gov/news/alerts/aa20-209a>
- [77] Sauvik Das, Laura A Dabbish, and Jason I Hong. 2019. A typology of perceived triggers for end-user security and privacy behaviors. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*.
- [78] Sauvik Das, Tiffany Hyun-Jin Kim, Laura A Dabbish, and Jason I Hong. 2014. The effect of social influence on security sensitivity. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 143–157.
- [79] Emma Davis. [n. d.]. QNAP Tells How To Deal With QSnatch — How To Fix Guide. <https://howtofix.guide/qnap-tells-how-to-deal-with-qsnatch/>
- [80] Hans De Bruijn and Ernst Ten Heuvelhof. 2018. *Management in networks*. Routledge.
- [81] Michele De Donno, Nicola Dragoni, Alberto Giarretta, and Angelo Spognardi. 2017. Analysis of DDoS-capable IoT malwares. In *Proceedings of the 2017 Federated Conference on Computer Science and Information Systems, FedCSIS 2017*. <https://doi.org/10.15439/2017F288>
- [82] Michele De Donno, Nicola Dragoni, Alberto Giarretta, and Angelo Spognardi. 2018. DDoS-capable IoT malwares: Comparative analysis and Mirai investigation. *Security and Communication Networks* 2018 (2018).
- [83] Charles DeBeck, Joshua Chung, and Dave McMillen. [n. d.]. I Can’t Believe Mirais: Tracking the Infamous IoT Malware. <https://securityintelligence.com/posts/i-cant-believe-mirais-tracking-the-infamous-iot-malware-2/>

- [84] Ozgur Dedehayir, Roland J Ortt, Carla Riverola, and Francesc Miralles. 2020. Innovators and early adopters in the diffusion of innovations: A literature review. *Digital Disruptive Innovation* (2020), 85–115.
- [85] Jyoti Deogirikar and Amarsinh Vidhate. 2017. Security attacks in IoT: A survey. In *2017 International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)(I-SMAC)*. IEEE, 32–37.
- [86] Department for Digital Culture Media & Sport. 2018. Code of Practice for consumer IoT security - GOV.UK. <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security/code-of-practice-for-consumer-iot-security>
- [87] Department for Digital Culture Media & Sport. 2018. Code of Practice for consumer IoT security - GOV.UK. <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security/code-of-practice-for-consumer-iot-security>
- [88] David Dittrich and Erin Kenneally. 2012. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. http://www.caيدا.org/publications/papers/2012/menlo_report_actual_formatted
- [89] David Dittrich, Erin Kenneally, et al. 2012. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. http://www.caيدا.org/publications/papers/2012/menlo_report_actual_formatted. Accessed:2021-05-25.
- [90] DNSFilter. [n. d.]. 2022 Threats by the Numbers. <https://www.dnsfilter.com/blog/2022-threats-by-the-numbers>
- [91] Steve Dodier-Lazaro, Ruba Abu-Salma, Ingolf Becker, and M Angela Sasse. 2017. From paternalistic to user-centred security: Putting users first with value-sensitive design. In *CHI 2017 Workshop on Values in Computing*. Values In Computing.
- [92] Michael Dodson, Daniel Thomas, and Alastair R Beresford. 2020. When will my PLC support Mirai? The security economics of large-scale attacks against Internet-connected ICS devices. (2020). <https://doi.org/10.17863/CAM.59520> ECRIME 2020 – SYMPOSIUM ON ELECTRONIC CRIME RESEARCH.
- [93] Domoticz. 2019. Domoticz Downloads. <https://www.domoticz.com/downloads/>

- [94] Paul Dourish, Rebecca E Grinter, Jessica Delgado De La Flor, and Melissa Joseph. 2004. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing* 8, 6 (2004), 391–401.
- [95] Zakir Durumeric, James Kasten, David Adrian, J. Alex Halderman, Michael Bailey, Frank Li, Nicholas Weaver, Johanna Amann, Jethro Beekman, Mathias Payer, and Vern Paxson. 2014. The matter of heartbleed. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC*. <https://doi.org/10.1145/2663716.2663755>
- [96] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, et al. 2014. The matter of Heartbleed. In *Proceedings of the 2014 conference on internet measurement conference*. 475–488.
- [97] Antoine d’Estalénx and Carlos Gañán. 2021. NURSE: eNd-UseR IoT malware detection tool for Smart homEs. In *Proceedings of the 11th International Conference on the Internet of Things (IoT ’21)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3494322.3494340>
- [98] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You’ve been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1065–1074.
- [99] Serge Egelman, David Molnar, Nicolas Christin, Alessandro Acquisti, Cormac Herley, and Shriram Krishnamurthi. 2010. Please continue to hold: An empirical study on user tolerance of security delays. In *Workshop on the Economics of Information Security (WEIS)*.
- [100] Serge Egelman and Stuart Schechter. 2013. The importance of being earnest [in security warnings]. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. https://doi.org/10.1007/978-3-642-39884-1_5
- [101] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the Experts: What Should Be on an IoT Privacy and Security Label?. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 447–464.
- [102] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In

- Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300764>
- [103] ENISA. 2018. Cybersecurity culture guidelines: behavioural aspects of cybersecurity. *European Union Agency for Network and Information Security* (2018).
- [104] ENISA. 2019. Good Practices for Security of IoT - Secure Software Development Lifecycle — ENISA. <https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot-1>
- [105] ENISA. 2023. Learn more about EU Cybersecurity Certification. <https://www.enisa.europa.eu/topics/certification/eu-cybersecurity-certification-faq>
- [106] European Commission. 2019. The Internet of Things | Shaping Europe’s digital future. <https://ec.europa.eu/digital-single-market/en/internet-of-things>
- [107] Eurofins. 2023. Connected Devices: Are they Secure and Fit for Purpose? <https://www.eurofins-cybersecurity.com/news/connected-devices-introduction/>
- [108] European Commission. 2020. The EU cybersecurity certification framework | Shaping Europe’s digital future. <https://ec.europa.eu/digital-single-market/en/eu-cybersecurity-certification-framework>
- [109] European Commission and the High Representative of the Union for Foreign Affairs and Security Policy. 2022. Study on Domain Name System (DNS) abuse - Publications Office of the EU. <https://op.europa.eu/en/publication-detail/-/publication/7d16c267-7f1f-11ec-8c40-01aa75ed71a1/language-en/>
- [110] European Telecommunications Standards Institute. 2020. ETSI - Consumer IoT security. <https://www.etsi.org/technologies/consumer-iot-security>
- [111] European Union & Agency for Network and Information Security. 2017. Baseline security recommendations for IoT in the context of critical information infrastructures - Publications Office of the EU. <https://op.europa.eu/en/publication-detail/-/publication/c37f8196-d96f-11e7-a506-01aa75ed71a1/language-en>. Accessed:2021-05-25.

- [112] European Union Agency for Cybersecurity (ENISA). 2018. *Cybersecurity Culture Guidelines: Behavioural Aspects of Cybersecurity*. <https://www.enisa.europa.eu/publications/cybersecurity-culture-guidelines-behavioural-aspects-of-cybersecurity>. Accessed:2021-05-25.
- [113] European Union and Council of Europe. 2004. *Document Library | Europass*. <https://europa.eu/europass/en/document-library>. Accessed: 2021-05-25.
- [114] Michael Fagan and Mohammad Maifi Hasan Khan. 2016. *Why Do They Do What They Do?: A Study of What Motivates Users to (Not) Follow Computer Security Advice*. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 59–75. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/fagan>
- [115] Michael Fagan, Katerina Megas, Karen Scarfone, and Matthew Smith. 2020. *Foundational Cybersecurity Activities for IoT Device Manufacturers*. <https://doi.org/10.6028/NIST.IR.8259>
- [116] Michael Fagan, Katerina N Megas, Karen Scarfone, and Matthew Smith. 2020. *IoT Device Cybersecurity Capability Core Baseline*. <https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259A.pdf>
- [117] Michael Fagan, Mary Yang, Allen Tan, Lora Randolph, and Karen Scarfone. 2019. *Security Review of Consumer Home Internet of Things (IoT) Products*. <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8267-draft.pdf>.
- [118] Federal Trade Commission. 2016. *ASUS Settles FTC Charges That Insecure Home Routers and “Cloud” Services Put Consumers’ Privacy At Risk | Federal Trade Commission*. <https://www.ftc.gov/news-events/press-releases/2016/02/asus-settles-ftc-charges-insecure-home-routers-cloud-services-p>
[ut?utm{ }source=govdelivery](https://www.ftc.gov/news-events/press-releases/2016/02/asus-settles-ftc-charges-insecure-home-routers-cloud-services-p)
- [119] Mohamed Ali Feki, Fahim Kawsar, Mathieu Boussard, and Lieven Trappeniers. 2013. *The internet of things: the next technological revolution*. *Computer 2* (2013), 24–25.
- [120] Adrienne Porter Felt, Alex Ainslie, Robert W. Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettis, Helen Harris, and Jeff Grimes. 2015. *Improving SSL warnings: Comprehension and adherence*. In *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/2702123.2702442>

- [121] Xuan Feng, Qiang Li, Haining Wang, and Limin Sun. 2018. Acquisitional Rule-based Engine for Discovering Internet-of-Things Devices. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 327–341. <https://www.usenix.org/conference/usenixsecurity18/presentation/feng>
- [122] Karen Fitzner and Elizabeth Heckinger. 2010. Sample size calculation and power analysis: a quick review. *The Diabetes Educator* 36, 5 (2010), 701–707.
- [123] Brian J Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*. 1–7.
- [124] Brian J Fogg. 2019. *Tiny habits: The small changes that change everything*. Eamon Dolan Books.
- [125] Brian J Fogg and Jason Hreha. 2010. Behavior wizard: A method for matching target behaviors with solutions. In *International Conference on Persuasive Technology*. Springer, 117–131.
- [126] Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang. 2016. Do or do not, there is no try: user engagement may not improve security outcomes. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. 97–111.
- [127] Alisa Frik, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. 2019. A promise is a promise: The effect of commitment devices on computer security intentions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [128] Steven Furnell. 2007. Making security usable: Are things improving? *Computers & Security* 26, 6 (2007), 434–443.
- [129] G. Huston. 2016. Measuring the End User. <https://labs.apnic.net/presentations/store/2016-02-10-ad-measurement.pdf>
- [130] Mario Galluscio, Nataliia Neshenko, Elias Bou-Harb, Yongliang Huang, Nasir Ghani, Jorge Crichigno, and Georges Kaddoum. 2017. A first empirical look on internet-scale exploitations of IoT devices. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE, 1–7. <https://doi.org/10.1109/PIMRC.2017.8292628>

- [131] Mario Galluscio, Nataliia Neshenko, Elias Bou-Harb, Yongliang Huang, Nasir Ghani, Jorge Crichigno, and Georges Kaddoum. 2017. A first empirical look on internet-scale exploitations of IoT devices. In *2017 IEEE 28th annual international symposium on personal, indoor, and mobile radio communications (PIMRC)*. IEEE, 1–7.
- [132] GCA. 2019. GCA | Global Cyber Alliance | Working to Eradicate Cyber Risk. <https://www.globalcyberalliance.org>
- [133] Christine Geeng and Franziska Roesner. 2019. Who’s In Control? Interactions In Multi-User Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [134] Michaelson George. 2012. Measuring the Internet for fun and profit. <https://labs.apnic.net/?p=83>
- [135] Google. 2009. Official Google Blog: Introducing Google Public DNS. <https://googleblog.blogspot.com/2009/12/introducing-google-public-dns.html>
- [136] Harm Griffioen and Christian Doerr. 2020. Examining Mirai’s Battle over the Internet of Things. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 743–756.
- [137] Sarthak Grover and Nick Feamster. 2016. The internet of unpatched things. *Proc. FTC PrivacyCon* (2016).
- [138] Hang Guo and John Heidemann. 2020. IoTSTEED: Bot-side Defense to IoT-based DDoS Attacks (Extended). *USC/Information Sciences Institute, Tech. Rep. ISI-TR-738* (2020).
- [139] Julie Haney, Yasemin Acar, and Susanne Furman. 2021. “It’s the Company, the Government, You and I”: User Perceptions of Responsibility for Smart Home Privacy and Security. In *30th USENIX Security Symposium (USENIX Security 21)*.
- [140] Julie M Haney, Yasemin Acar, and Susanne M Furman. 2021. "It’s the Company, the Government, You and I": User Perceptions of Responsibility for Smart Home Privacy and Security. In *30th USENIX Security Symposium (USENIX Security ‘21)*. USENIX Association, Vancouver, B.C. <https://www.usenix.org/conference/usenixsecurity21/presentation/haney>
- [141] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. 2019. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium (USENIX Security ‘19)*. 105–122.

- [142] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlence Fernandes, and Blase Ur. 2018. Rethinking access control and authentication for the home internet of things (IoT). In *27th USENIX Security Symposium (USENIX Security '18)*. 255–272.
- [143] Weijia He, Jesse Martinez, Roshni Padhi, Lefan Zhang, and Blase Ur. 2019. When smart devices are stupid: negative experiences using home smart devices. In *2019 IEEE Security and Privacy Workshops (SPW)*. IEEE, 150–155.
- [144] Hikvision. 2021. DS-2TD1117-2/PA | HeatPro Series | Hikvision. <https://www.hikvision.com/en/products/Thermal-Products/Security-thermal-cameras/heatpro-series/ds-2td1117-2-pa/>
- [145] HP. 2014. HP News - HP Study Reveals 70 Percent of Internet of Things Devices Vulnerable to Attack. <https://www8.hp.com/us/en/hp-news/press-release.html?id=1744676>
- [146] Paul Hünermund and Beyers Louw. 2020. On the nuisance of control variables in regression analysis. *arXiv preprint arXiv:2005.10314* (2020).
- [147] Geoff Huston. 2022. Recursive resolvers. <https://stats.labs.apnic.net/rvr.csv/>
- [148] IETF. 2023. IETF | Introduction. <https://www.ietf.org/about/introduction/>
- [149] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2021. LZR: Identifying Unexpected Internet Services. In *30th USENIX Security Symposium*.
- [150] M James. 2017. Secure by Design: Improving the cyber security of consumer Internet of Things Report. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/775559/Secure_by_Design_Report_.pdf
- [151] Talha Javed, Muhammad Haseeb, Muhammad Abdullah, and Mobin Javed. 2020. Using Application Layer Banner Data to Automatically Identify IoT Devices. *SIG-COMM Comput. Commun. Rev.* 50, 3 (July 2020), 23–29. <https://doi.org/10.1145/3411740.3411744>
- [152] Mohammad Hanif Jhaveri, Orcun Cetin, Carlos Gañán, Tyler Moore, and Michel Van Eeten. 2017. Abuse reporting and the fight against cybercrime. *Comput. Surveys* (2017). <https://doi.org/10.1145/3003147>

- [153] Buchner Johannes. [n. d.]. ImageHash · PyPI. <https://pypi.org/project/ImageHash/>
- [154] Bonnie E John and Steven J Marks. 1997. Tracking the effectiveness of usability evaluation methods. *Behaviour & Information Technology* 16, 4-5 (1997), 188–202.
- [155] Shane D Johnson, John M Blythe, Matthew Manning, and Gabriel TW Wong. 2020. The impact of IoT security labelling on consumer product choice and willingness to pay. *PloS one* 15, 1 (2020), e0227800.
- [156] G. Kambourakis, C. Kolias, and A. Stavrou. 2017. The Mirai botnet and the IoT Zombie Armies. In *MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM)*. 267–272. <https://doi.org/10.1109/MILCOM.2017.8170867>
- [157] Edward L Kaplan and Paul Meier. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53, 282 (1958), 457–481.
- [158] Kaspersky. 2019. New Mirai botnet is targeting enterprise IoT | Kaspersky official blog. <https://www.kaspersky.com/blog/mirai-enterprise/26032/>
- [159] Erin Kenneally. 2019. Economics and Incentives Driving IoT Privacy and Security, Pt. 1. *IEEE Internet of Things Magazine* 2, 1 (2019), 6–7.
- [160] John P Klein and Melvin L Moeschberger. 2003. *Survival analysis: techniques for censored and truncated data*. Vol. 1230. Springer.
- [161] David G Kleinbaum, K Dietz, M Gail, Mitchel Klein, and Mitchell Klein. 2002. *Logistic regression*. Springer.
- [162] Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. 2017. DDoS in the IoT: Mirai and other botnets. *Computer* (2017). <https://doi.org/10.1109/MC.2017.201>
- [163] David Kotz and Travis Peters. 2017. Challenges to ensuring human safety throughout the life-cycle of Smart Environments. In *Proceedings of the 1st ACM Workshop on the Internet of Safe Things*. 1–7.
- [164] Kat Krol, Matthew Moroz, and M. Angela Sasse. 2012. Don’t work. Can’t work? Why it’s time to rethink security warnings. In *7th International Conference on Risks and Security of Internet and Systems, CRiSIS 2012*. <https://doi.org/10.1109/CRISIS.2012.6378951>

- [165] Kat Krol, Matthew Moroz, and M Angela Sasse. 2012. Don't work. Can't work? Why it's time to rethink security warnings. In *2012 7th international conference on risks and security of internet and systems (CRiSIS)*. IEEE, 1–8.
- [166] Deepak Kumar, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. 2019. All Things Considered: An Analysis of IoT Devices on Home Networks. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 1169–1185. <https://www.usenix.org/conference/usenixsecurity19/presentation/kumar-deepak>
- [167] E. Lear, R. Droms, and D. Romascanu. 2019. *Manufacturer Usage Description Specification*. Technical Report. <https://doi.org/10.17487/RFC8520>
- [168] Scott T Leatherdale. 2019. Natural experiment methodology for research: a review of how different methods can support real-world research. *International Journal of Social Research Methodology* 22, 1 (2019), 19–35.
- [169] Gwanhoo Lee. 2019. What roles should the government play in fostering the advancement of the Internet of Things? *Telecommunications Policy* 43, 5 (2019), 434–444.
- [170] Eireann Leverett, Richard Clayton, and Ross Anderson. 2017. Standardisation and Certification of the 'Internet of Things'. In *Proceedings of WEIS*. 1–24.
- [171] Clayton Lewis. 1982. *Using the "thinking-aloud" method in cognitive interface design*. IBM TJ Watson Research Center Yorktown Heights, NY.
- [172] Frank Li, Michael Bailey, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Damon McCoy, Stefan Savage, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *USENIX Security Symposium*.
- [173] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating web hijacking: Notification effectiveness and webmaster comprehension. In *25th International World Wide Web Conference, WWW 2016*. <https://doi.org/10.1145/2872427.2883039>
- [174] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating web hijacking: Notification effectiveness and

- webmaster comprehension. In *Proceedings of the 25th International Conference on World Wide Web*. 1009–1019.
- [175] Ming Liu, Zhi Xue, Xiangjian He, and Jinjun Chen. 2019. Cyberthreat-Intelligence Information Sharing: Enhancing Collaborative Security. *IEEE Consumer Electronics Magazine* 8, 3 (may 2019), 17–22. <https://doi.org/10.1109/MCE.2019.2892220>
- [176] Jason Livingood, Nirmal Mody, and Mike O’Reirdan. 2012. Recommendations for the Remediation of Bots in ISP Networks. <https://tools.ietf.org/html/rfc6561>
- [177] Qasim Lone, Maciej Korczyński, Carlos Gañán, and Michel van Eeten. 2020. SAVING the Internet: Explaining the Adoption of Source Address Validation by Internet Service Providers. In *Workshop on the Economics of Information Security*.
- [178] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. 2021. Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 2489–2506. <https://www.usenix.org/conference/usenixsecurity21/presentation/maass>
- [179] Malware Wiki. 2019. Mirai | Malware Wiki l. <https://malware.wikia.org/wiki/Mirai>
- [180] Emeline Marechal and Benoît Donnet. 2020. Network Fingerprinting: Routers under Attack. In *IEEE International Workshop on Traffic Measurements for Cybersecurity (WTMC)*. <https://orbi.uliege.be/handle/2268/248733>
- [181] J. Margolis, T. T. Oh, S. Jadhav, Y. H. Kim, and J. N. Kim. 2017. An In-Depth Analysis of the Mirai Botnet. In *2017 International Conference on Software Security and Assurance (ICSSA)*. 6–12. <https://doi.org/10.1109/ICSSA.2017.12>
- [182] Anna Marton. 2021. Protecting Consumers of IoT Products from Cyber Threats - Latest Legislation from around the Globe Explained - IoTAC Insights. <https://iotac.eu/protecting-consumers-of-iot-products-from-cyber-threats-latest-legislation-from-around-the-globe/>
- [183] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.

- [184] Louis Meuleman. 2008. *Public management and the metagovernance of hierarchies, networks and markets: The feasibility of designing and managing governance style combinations*. Springer Science & Business Media.
- [185] Susan Michie, Maartje M Van Stralen, and Robert West. 2011. The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation science* 6, 1 (2011), 1–12.
- [186] Vladimir Kuskov Mikhail Kuzin, Yaroslav Shmelev. 2018. New trends in the world of IoT threats | Securelist. <https://securelist.com/new-trends-in-the-world-of-iot-threats/87991/>
- [187] Ministry of Economic Affairs and Climate Policy. [n. d.]. Roadmap for Digital Hard- and Software Security | Report | Government.nl. <https://www.government.nl/documents/reports/2018/04/02/roadmap-for-digital-hard--and-softw-are-security>
- [188] Ministry of Economic Affairs and Climate Policy. 2020. The Radiocommunications Agency pushes for eight security requirements for smart devices | Dutch Authority for Digital Infrastructure | Rijksinspectie Digitale Infrastructuur (RDI). <https://www.rdi.nl/radiocommunications-agency/news/2020/08/26>
- [189] Netherlands Chamber of Commerce Ministry of Economic Affairs and Climate Policy and Netherlands Enterprise Agency. 2022. New rules for smart devices and digital products | Business.gov.nl. <https://business.gov.nl/running-your-business/business-management/legal-matters/new-rules-for-smart-devices-and-digital-products/>
- [190] David Moore. 2002. Network Telescopes: Observing Small or Distant Security Events. In *11th USENIX Security Symposium (USENIX Security '02)*. USENIX Association, San Francisco, CA. <https://www.usenix.org/conference/11th-usenix-security-symposium/network-telescopes-observing-small-or-distant-security>
- [191] Philipp Morgner, Christoph Mai, Nicole Koschate-Fischer, Felix Freiling, and Zinaida Benenson. 2020. Security update labels: Establishing economic incentives for security patching of IoT consumer products. In *2020 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 429–446.
- [192] P. Morgner, C. Mai, N. Koschate-Fischer, F. Freiling, and Z. Benenson. 2020. Security Update Labels: Establishing Economic Incentives for Security Patching of

- IoT Consumer Products. In *2020 IEEE Symposium on Security and Privacy (SP)*. 429–446.
- [193] Mozilla. 2021. *Privacy not included. <https://foundation.mozilla.org/en/privacy/privacy-not-included/>. Accessed: 2021-05-25.
- [194] Asuka Nakajima, Takuya Watanabe, Eitaro Shioji, Mitsuaki Akiyama, and Maverick Woo. 2019. A pilot study on consumer IoT device vulnerability disclosure and patch release in Japan and the United States. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*. 485–492.
- [195] Iain Nash. 2020. A Proposed Civil Liability Framework for Disrupting Botnets, with a particular focus on Smart Devices. (2020). <https://www.youtube.com/watch?v=NshPs3GxRcA> Botconf.
- [196] National Cybersecurity Centre. 2019. Smart devices: using them safely in your home. <https://www.ncsc.gov.uk/guidance/smart-devices-in-the-home>
- [197] National Cybersecurity Centre. 2020. Protective Domain Name Service (PDNS) - NCSC.GOV.UK. <https://www.ncsc.gov.uk/information/pdns>
- [198] National Initiative for Cybersecurity Careers. 2019. DHS NCSAM 2019 - 5 Steps Protecting Your Digital Home. <https://www.cisa.gov/sites/default/files/publications/Five-Steps-to-Protecting-Your-Digital-Home-Tip-Sheet-122019-508.pdf>
- [199] National Security Agency. 2021. *Selecting a Protective DNS Service Why Protective DNS? How does it work?* Technical Report.
- [200] Nataliia Neshenko, Elias Bou-Harb, Jorge Crichigno, Georges Kaddoum, and Nasir Ghani. 2019. Demystifying IoT security: an exhaustive survey on IoT vulnerabilities and a first empirical look on internet-scale IoT exploitations. *IEEE Communications Surveys & Tutorials* 21, 3 (2019), 2702–2733.
- [201] Nataliia Neshenko, Martin Husak, Elias Bou-Harb, Pavel Celeda, Sameera Al-Mulla, and Claude Fachkha. 2019. Data-Driven Intelligence for Characterizing Internet-Scale IoT Exploitations. In *2018 IEEE Globecom Workshops, GC Wkshps 2018 - Proceedings*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/GLOCOMW.2018.8644468>

- [202] Netscout. 2018. Dawn of the terrorbit era. https://www.netscout.com/sites/default/files/2019-02/SECR_001_EN-1901%20-%20NETSCOUT%20Threat%20Intelligence%20Report%20H%202018.pdf. Accessed: 2021-05-25.
- [203] IoT Business News. 2020. IoT News - The IoT in 2030: 24 billion connected things generating \$1.5 trillion - IoT Business News. <https://iotbusinessnews.com/2020/05/20/03177-the-iot-in-2030-24-billion-connected-things-generating-1-5-trillion/>
- [204] Quoc-Dung Ngo, Huy-Trung Nguyen, Van-Hoang Le, and Doan-Hieu Nguyen. 2020. A survey of IoT malware and detection methods based on static features. *ICT Express* 6, 4 (2020), 280–286.
- [205] Kenneth D Nguyen, Heather Rosoff, and Richard S John. 2017. Valuing information security from a phishing attack. *Journal of Cybersecurity* 3, 3 (2017), 159–171.
- [206] Larissa Nicholls, Yolande Strengers, and Jathan Sadowski. 2020. Social impacts and control in the smart home. *Nature Energy* 5, 3 (2020), 180–182.
- [207] Jakob Nielsen. 1994. Enhancing the explanatory power of usability heuristics. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 152–158.
- [208] Alexandra Nisenoff, Nick Feamster, Madeleine A Hoofnagle, and Sydney Zink. 2021. User Expectations and Understanding of Encrypted DNS Settings. In *Proc. NDSS DNS Privacy Workshop, Virtual Event*.
- [209] Alexandra Nisenoff, Ranya Sharma, and Nick Feamster. 2022. Understanding User Awareness and Behaviors Concerning Encrypted DNS Settings. *arXiv preprint arXiv:2208.04991* (2022).
- [210] Arman Noroozian, Michael Ciere, Maciej Korczynski, Samaneh Tajalizadehkhoob, and Michel Van Eeten. 2017. Inferring the Security Performance of Providers from Noisy and Heterogenous Abuse Datasets. In *16th Workshop on the Economics of Information Security*. http://weis2017.econinfosec.org/wp-content/uploads/sites/3/2017/05/WEIS_2017_paper_60.pdf.
- [211] NOS. 2019. Kabinet wil verplichte updates voor 'slimme' apparaten | NOS. <https://nos.nl/artikel/2315967-kabinet-wil-verplichte-updates-voor-slimme-apparaten.html>

- [212] OECD. 2019. Roadmap: Measuring the Internet of Things. In *Measuring the Digital Transformation*. OECD, 108–109. <https://doi.org/10.1787/2abc4f98-en>
- [213] Organisation for Economic Co-operation and Development. 2010. The Economic and Social Role of Internet Intermediaries. *Notes* April (2010), 49. <http://www.oecd.org/dataoecd/49/4/44949023.pdf>
- [214] Magda Osman, Scott McLachlan, Norman Fenton, Martin Neil, Ragnar Löfstedt, and Björn Meder. 2020. Learning from behavioural changes that fail. *Trends in Cognitive Sciences* (2020).
- [215] OWASP Foundation. 2018. OWASP Internet of Things Project - OWASP. https://wiki.owasp.org/index.php/OWASP_Internet_of_Things_Project#tab=IoT_Top_10
- [216] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. 2015. IoTPOT: Analysing the Rise of IoT Compromises. In *9th {USENIX} Workshop on Offensive Technologies ({WOOT} 15)*. {USENIX} Association, Washington, D.C. <https://www.usenix.org/conference/woot15/workshop-program/presentation/pa>
- [217] Simon Parkin, Elissa M Redmiles, Lynne Coventry, and M Angela Sasse. 2019. Security when it is welcome: Exploring device purchase as an opportune moment for security behavior change. In *Proceedings of the Workshop on Usable Security and Privacy (USEC'19)*. Internet Society.
- [218] Roberto Perdisci, Thomas Papastergiou, Omar Alrawi, and Manos Antonakakis. 2020. IoTFinder : Efficient Large-Scale Identification of IoT Devices via Passive DNS Traffic Analysis. *Acm Imc* (2020).
- [219] Karine Perset. 2010. The economic and social role of Internet intermediaries. (2010).
- [220] Yusuf Perwej, Kashiful Haq, Firoj Parwej, M Mumdouh, and Mohamed Hassan. 2019. The internet of things (IoT) and its application domains. *International Journal of Computer Applications* 182, 49 (2019), 36–49.
- [221] Shari Lawrence Pfleeger, M Angela Sasse, and Adrian Furnham. 2014. From weakest link to security hero: Transforming staff security behavior. *Journal of Homeland Security and Emergency Management* 11, 4 (2014), 489–510.

- [222] Jeroen Pijpker and Harald Vranken. 2016. The Role of Internet Service Providers in Botnet Mitigation. In *2016 European Intelligence and Security Informatics Conference (EISIC)*. IEEE, 24–31. <https://doi.org/10.1109/EISIC.2016.013>
- [223] Erika Shehan Poole, Marshini Chetty, Tom Morgan, Rebecca E Grinter, and W Keith Edwards. 2009. Computer help at home: methods and motivations for informal technical support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 739–748.
- [224] James O Prochaska, Colleen A Redding, Kerry E Evers, et al. 2015. The transtheoretical model and stages of change. *Health behavior: Theory, research, and practice* 97 (2015).
- [225] Prolific. 2022. Prolific · Quickly find research participants you can trust. <https://www.prolific.co/>
- [226] QNAP. 2020. How to prevent attacks by QSnatch | QNAP. <https://www.qnap.com/en/how-to/knowledge-base/article/about-qsnatch>
- [227] QNAP. 2020. Security Advisory for Malware QSnatch - Security Advisory | QNAP. <https://www.qnap.com/en/security-advisory/nas-201911-01>
- [228] QNAP. 2021. Security Advisory for Malware QSnatch - Security Advisory | QNAP. <https://www.qnap.com/en/security-advisory/nas-201911-01>. Accessed: 2021-05-25.
- [229] Emilee Rader, Rick Wash, and Brandon Brooks. 2012. Stories as informal lessons about security. In *Proceedings of the Eighth Symposium on Usable Privacy and Security - SOUPS '12*. ACM Press, New York, New York, USA, 1. <https://doi.org/10.1145/2335356.2335364>
- [230] Roxana Radu and Michael Hausding. 2020. Consolidation in the DNS resolver market—how much, how fast, how dangerous? *Journal of Cyber Policy* 5, 1 (2020), 46–64.
- [231] Elissa M Redmiles. 2019. "Should I Worry?" A Cross-Cultural Examination of Account Security Incident Response. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 920–934.
- [232] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2017. Where is the digital divide? a survey of security, privacy, and socioeconomics. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 931–936.

- [233] Elissa M Redmiles, Amelia R Malone, and Michelle L Mazurek. 2016. I think they're trying to tell me something: Advice sources and selection for digital security. In *2016 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 272–288.
- [234] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. 2020. A Comprehensive Quality Evaluation of Security and Privacy Advice on the Web. In *29th USENIX Security Symposium (USENIX Security '20)*. USENIX Association, 89–108. <https://www.usenix.org/conference/usenixsecurity20/presentation/redmiles>
- [235] Elissa M Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L Mazurek. 2020. A comprehensive quality evaluation of security and privacy advice on the web. In *29th USENIX Security Symposium (USENIX Security 20)*. 89–108.
- [236] Elissa M Redmiles, Ziyun Zhu, Sean Kross, Dhruv Kuchhal, Tudor Dumitras, and Michelle L Mazurek. 2018. Asking for a friend: Evaluating response biases in security user studies. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 1238–1255.
- [237] Robert W Reeder, Iulia Ion, and Sunny Consolvo. 2017. 152 simple steps to stay safe online: Security advice for non-tech-savvy users. *IEEE Security & Privacy* 15, 5 (2017), 55–64.
- [238] Karen Renaud and Wendy Goucher. 2014. The curious incidence of security breaches by knowledgeable employees and the pivotal role a of security culture. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*. Springer, 361–372.
- [239] Jason T Rich, J Gail Neely, Randal C Paniello, Courtney CJ Voelker, Brian Nussenbaum, and Eric W Wang. 2010. A practical guide to understanding Kaplan-Meier curves. *Otolaryngology—Head and Neck Surgery* 143, 3 (2010), 331–336.
- [240] Michael Richardson and M Ranganathan. 2019. *Manufacturer Usage Description for quarantined access to firmware*. Technical Report draft-richardson-shg-mud-quarantined-access-01. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-richardson-shg-mud-quarantined-access-01>
- [241] Horst WJ Rittel and Melvin M Webber. 1973. Dilemmas in a general theory of planning. *Policy sciences* 4, 2 (1973), 155–169.

- [242] Robert Graham. [n. d.]. GitHub - robertdavidgraham/masscan: TCP port scanner, spews SYN packets asynchronously, scanning entire Internet in under 5 minutes. <https://github.com/robertdavidgraham/masscan>
- [243] Elsa Rodríguez, Arman Noroozian, Michel van Eeten, and Carlos Gañán. 2021. Superspreaders: Quantifying the Role of IoT Manufacturers in Device Infections. *Workshop on the Economics of Information Security (WEIS)* (2021).
- [244] Elsa Rodríguez, Susanne Verstegen, Arman Noroozian, Daisuke Inoue, Takahiro Kasama, Michel van Eeten, and Carlos H Gañán. 2021. User compliance and remediation success after IoT malware notifications. *Journal of Cybersecurity* 7, 1 (2021).
- [245] Sabine Roeser. 2011. Nuclear energy, risk, and emotions. *Philosophy & Technology* 24 (2011), 197–201.
- [246] Brent Rowe and Dallas Wood. 2013. Are home internet users willing to pay ISPs for improvements in cyber security? In *Economics of information security and privacy III*. Springer, 193–212.
- [247] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Daniel J. Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. 2020. A Haystack Full of Needles: Scalable Detection of IoT Devices in the Wild. *Acm Imc* (2020). arXiv:arXiv:2009.01880v1
- [248] Angela Sasse. 2015. Scaring and bullying people into security won’t work. *IEEE Security & Privacy* 13, 3 (2015), 80–83.
- [249] Bruce Schneier. 2014. The internet of things is wildly insecure-and often unpatchable. *Schneier on Security* 6 (2014).
- [250] Bruce Schneier. 2018. *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company.
- [251] William Seymour, Martin J Kraemer, Reuben Binns, and Max Van Kleek. 2020. Informing the design of privacy-empowering tools for the connected home. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [252] Shadowserver. 2019. Drone/Botnet-Drone Report | Shadowserver. <https://www.shadowserver.org/what-we-do/network-reporting/drone-botnet-drone-report/>

- [253] Shadowserver. 2021. Drone/Botnet-Drone Report | Shadowserver. <https://www.shadowserver.org/what-we-do/network-reporting/drone-botnet-drone-report/>. Accessed: 2021-05-25.
- [254] Shadowserver. 2021. Index of Iot Exposed Infected Device Stats. <https://cra.ci.rcl.lu/opendata/variote/iot-exposed-infected-device-stats/>
- [255] Shadowserver. 2022. The Shadowserver Foundation. <https://www.shadowserver.org/>
- [256] Paschal Sheeran and Thomas L Webb. 2016. The intention–behavior gap. *Social and personality psychology compass* 10, 9 (2016), 503–518.
- [257] Shodan. 2019. Shodan. <https://www.shodan.io/>
- [258] Shodan. 2020. . <https://www.shodan.io/>
- [259] H. Sinanović and S. Mrdovic. 2017. Analysis of Mirai malicious software. In *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. 1–5. <https://doi.org/10.23919/SOFTCOM.2017.8115504>
- [260] Arunan Sivanathan, Hassan Habibi Gharakheili, and Vijay Sivaraman. 2018. Can We Classify an IoT Device using TCP Port Scan?. In *2018 IEEE 9th International Conference on Information and Automation for Sustainability, ICIAfS 2018*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICIAfS.2018.8913346>
- [261] Michael Smithson and Jay Verkuilen. 2006. A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods* (2006). <https://doi.org/10.1037/1082-989X.11.1.54>
- [262] Software Tested. 2020. How to Get Rid of the Qsnatch Malware - Software Tested. <https://softwaretested.com/anti-malware/how-to-get-rid-of-the-qsnatch-malware/>
- [263] Yubo Song, Qiang Huang, Junjie Yang, Ming Fan, Aiqun Hu, and Yu Jiang. 2019. IoT Device Fingerprinting for Relieving Pressure in the Access Control. In *Proceedings of the ACM Turing Celebration Conference - China (ACM TURC '19)*. Association for Computing Machinery, New York, NY, USA, Article 143, 8 pages. <https://doi.org/10.1145/3321408.3326671>

- [264] Statista. 2023. IoT connected devices worldwide 2019-2030 | Statista. <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>
- [265] Jan-Benedict EM Steenkamp, Martijn G De Jong, and Hans Baumgartner. 2010. Socially desirable response tendencies in survey research. *Journal of Marketing Research* 47, 2 (2010), 199–214.
- [266] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. 2018. Didn't You Hear Me? – Towards More Successful Web Vulnerability Notifications. In *Proceedings of the 25th Annual Symposium on Network and Distributed System Security (NDSS '18)*. <https://publications.cispa.saarland/1190/>
- [267] Heshan Sun and Ping Zhang. 2006. The role of moderating factors in user technology acceptance. *International journal of human-computer studies* 64, 2 (2006), 53–78.
- [268] Symantec. 2019. *Internet Security Threat Report Volume 24* l. Technical Report.
- [269] Magdalena Szumilas. 2010. Explaining odds ratios. *Journal of the Canadian Academy of Child and Adolescent Psychiatry* 19, 3 (aug 2010), 227–229. <http://www.csm-oxford.org.uk/>
- [270] Leonie Tanczer, Irina Brass, Miles Elsdén, Madeline Carr, and Jason J Blackstock. 2019. The United Kingdom's emerging internet of things (IoT) policy landscape. Tanczer, LM, Brass, I., Elsdén, M., Carr, M., & Blackstock, J.(2019). *The United Kingdom's Emerging Internet of Things (IoT) Policy Landscape*. In R. Ellis & V. Mohan (Eds.), *Rewired: Cybersecurity Governance* (2019), 37–56.
- [271] Fumiyuki Tanemo, Mitsuhiro Osaki, Hiroaki Waki, Yutaka Ishioka, and Kazuhito Matsushita. 2020. A Method of Creating Data for Device-information Extraction by Efficient Wide-area-network Scanning of IoT Devices. In *2020 International Conference on Information Networking (ICOIN)*. IEEE, 643–648.
- [272] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 367–385.
- [273] The Mozilla Foundation. 2020. Mozilla - *privacy not included. <https://foundation.mozilla.org/en/privacynotincluded/>

- [274] Tung Liam. 2017. IoT Devices will outnumber the world's population this year for the first time | ZDNet. <https://www.zdnet.com/article/iot-devices-will-outnumber-the-worlds-population-this-year-for-the-first-time/>
- [275] Nastasha Tupas. 2021. ACSC launches new cyber guard for government data - Defence Connect. <https://www.defenceconnect.com.au/intel-cyber/8911-acsc-launches-new-cyber-guard-for-government-data>
- [276] Alex Turing, Hui Wang, and Liu Yang. 2021. New Threat: Matryosh Botnet Is Spreading. <https://blog.netlab.360.com/matryosh-botnet-is-spreading-en/>. Accessed: 2021-05-25.
- [277] UK Department for Digital, Culture, Media & Sport (DCMS). 2018. Code of Practice for Consumer IoT Security. <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security>.
- [278] Maaïke Van Den Haak, Menno De Jong, and Peter Jan Schellens. 2003. Retrospective vs. concurrent think-aloud protocols: testing the usability of an online library catalogue. *Behaviour & information technology* 22, 5 (2003), 339–351.
- [279] Wendelien Van Eerde. 2000. Procrastination: Self-regulation in initiating aversive goals. *Applied Psychology* 49, 3 (2000), 372–389.
- [280] Kami Vaniea, Emilee Rader, and Rick Wash. 2014. Betrayed by updates: How negative experiences affect future security. In *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/2556288.2557275>
- [281] Kami Vaniea and Yasmeen Rashidi. 2016. Tales of software updates: The process of updating software. In *Proceedings of the 2016 chi conference on human factors in computing systems*. 3215–3226.
- [282] Marie Vasek and Tyler Moore. 2012. Do Malware Reports Expedite Cleanup? An Experimental Study. In *Presented as part of the 5th Workshop on Cyber Security Experimentation and Test*. USENIX, Bellevue, WA. <https://www.usenix.org/conference/cset12/workshop-program/presentation/Vasek>
- [283] Marie Vasek and Tyler Moore. 2012. Do Malware Reports Expedite Cleanup? An Experimental Study.. In *5th Workshop on Cyber Security Experimentation and Test (CSET '12)*.

- [284] Benjamin Vignau, Raphaël Khoury, and Sylvain Hallé. 2019. 10 years of IoT malware: A feature-based taxonomy. In *2019 IEEE 19th International Conference on Software Quality, Reliability and Security Companion (QRS-C)*. IEEE, 458–465.
- [285] Aidan R Vining and David L Weimer. 1988. Information asymmetry favoring sellers: a policy framework. *Policy Sciences* 21, 4 (1988), 281–303.
- [286] Natalija Vlajic and Daiwei Zhou. 2018. IoT as a Land of Opportunity for DDoS Hackers. *Computer* (2018). <https://doi.org/10.1109/MC.2018.3011046>
- [287] Hui Wang. 2018. Fbot, A Satori Related Botnet Using Block-chain DNS System. <https://blog.netlab.360.com/threat-alert-a-new-worm-fbot-cleaning-adbminer-is-using-a-blockchain-based-dns-en/>.
- [288] Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security (SOUPS '10)*. Article 11, 16 pages.
- [289] Rick Wash and Emilee Rader. 2015. Too much knowledge? security beliefs and protective behaviors among united states internet users. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 309–325.
- [290] Rick Wash, Emilee Rader, Kami Vaniea, and Michelle Rizor. 2014. Out of the loop: How automated software updates cause unintended security consequences. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 89–104.
- [291] Sophia Waterfield. 2022. EU cyber resilience act: Europe aims for secure connected IoT devices. <https://techmonitor.ai/policy/privacy-and-data-protection/eu-cyber-resilience-act-iot-connected-devices>
- [292] Whalebone. 2022. Press Release: DNS4EU | Whalebone. <https://www.whalebone.io/post/press-release-dns4eu>
- [293] Michael S. Wogalter and Kenneth R. Laughery. 1996. Warning! Sign and label effectiveness. *Current Directions in Psychological Science* (1996). <https://doi.org/10.1111/1467-8721.ep10772712>
- [294] Chun-Jung Wu, Ying Tie, Satoshi Hara, Kazuki Tamiya, Akira Fujita, Katsunari Yoshioka, and Tsutomu Matsumoto. 2018. Iotprotect: Highly deployable whitelist-based protection for low-cost internet-of-things devices. *Journal of Information Processing* 26 (2018), 662–672.

- [295] Matheus Xavier Ferreira, S Matthew Weinberg, Danny Yuxing Huang, Nick Feamster, and Tithi Chattopadhyay. 2019. Selling a Single Item with Negative Externalities. In *The World Wide Web Conference*. 196–206.
- [296] Xuan Feng, Qiang Li, Haining Wang, and Limin Sun. 2016. Characterizing industrial control system devices on the Internet. In *2016 IEEE 24th International Conference on Network Protocols (ICNP)*. 1–10.
- [297] Yandex. 2022. Yandex DNS. <https://dns.yandex.com/>
- [298] Donghui Yang, Zhenyu Li, Haiyang Jiang, Gareth Tyson, Hongtao Li, Gaogang Xie, and Yu Zeng. 2022. A deep dive into DNS behavior and query failures. *Computer Networks* 214 (2022), 109131.
- [299] Kai Yang, Qiang Li, and Limin Sun. 2019. Towards automatic fingerprinting of IoT devices in the cyberspace. *Computer Networks* 148 (jan 2019), 318–327. <https://doi.org/10.1016/J.COMNET.2018.11.013>
- [300] Lihua Yin, Xi Luo, Chunsheng Zhu, Liming Wang, Zhen Xu, and Hui Lu. 2019. ConnSpooiler: Disrupting C&C Communication of IoT-Based Botnet through Fast Detection of Anomalous Domain Queries. *IEEE Transactions on Industrial Informatics* (2019), 1–1. <https://doi.org/10.1109/TII.2019.2940742>
- [301] Danny Yuxing Huang, Noah Apthorpe, Gunes Acar, Frank Li, and Nick Feamster. 2020. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT / Ubicomp)* (2020).
- [302] Liudmila Zavolokina, Manuel Schlegel, and Gerhard Schwabe. 2021. How can we reduce information asymmetries and enhance trust in ‘The Market for Lemons’? *Information Systems and e-Business Management* 19, 3 (2021), 883–908.
- [303] ZDNet. 2019. Thousands of QNAP NAS devices have been infected with the QSnatch malware | ZDNet. <https://www.zdnet.com/article/thousands-of-qnap-nas-devices-have-been-infected-with-the-qsnatch-malware/>. Accessed: 2021-05-25.
- [304] Eric Zeng, Shirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *thirteenth symposium on usable privacy and security (SOUPS 2017)*. 65–80.

- [305] Eric Zeng and Franziska Roesner. 2019. Understanding and improving security and privacy in multi-user smart homes: a design exploration and in-home user study. In *28th USENIX Security Symposium (USENIX Security '19)*. 159–176.
- [306] Verena Zimmermann, Paul Gerber, Karola Marky, Leon Böck, and Florian Kirchbuchner. 2019. Assessing Users' Privacy and Security Concerns of Smart Home Technologies. *i-com* 18, 3 (2019), 197–216.

APPENDIX A

A.1. CORRELATION BETWEEN THE STEPS PERFORMED BY CUSTOMERS

r_s	Step 1	Step 2	Step 3	Step 4	Step 5
Step 1	1				
Step 2	0.49	1			
Step 3	0.49	0.74	1		
Step 4	0.49	0.42	0.58	1	
Step 5	0.44	0.56	0.56	0.73	1

Figure 1: Correlation between the steps performed by customers.

A.2. LIKELIHOOD RATIO TEST COMPLIANCE MODELS

Likelihood ratio test Models 1-3

Model 1: compliance ratio ~ Walled Garden + Email-oly

Model 2: Compliance ratio ~ Walled Garden + Email-only + Age + Small business+ Male

Model 3: Compliance ratio ~ Walled Garden+ Email-only + Age + Small business+ Male + Domoticz

#Df LogLik Df Chisq Pr(>Chisq)

1 4 95.155

2 7 95.267 3 0.2240 0.9736

3 8 95.332 1 0.1319 0.7165

Likelihood ratio test Models 4-5

Model 1: Compliance ratio ~ Walled Garden + Age + Small business + Male + Domoticz + Understood notification

Model 2: Compliance ratio ~ Walled Garden + Age + Small business + Male + Domoticz + Understood notification+ Safe internet + Other motivation

#Df LogLik Df Chisq Pr(>Chisq)

4 8 15.630

5 10 22.185 2 13.11 0.001423 **

Signif. codes: 0 '****' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Note: Model 1-3 are not different with respect to likelihood value, and Model 5 shows improvement with respect to likelihood value of Model 4.

Table 1

A.3. SURVEY PROTOCOL

These are the survey protocols that were used to conduct the survey with the users in the different treatment groups. The survey was conducted in Dutch. We translated the questions as accurate as possible to English.

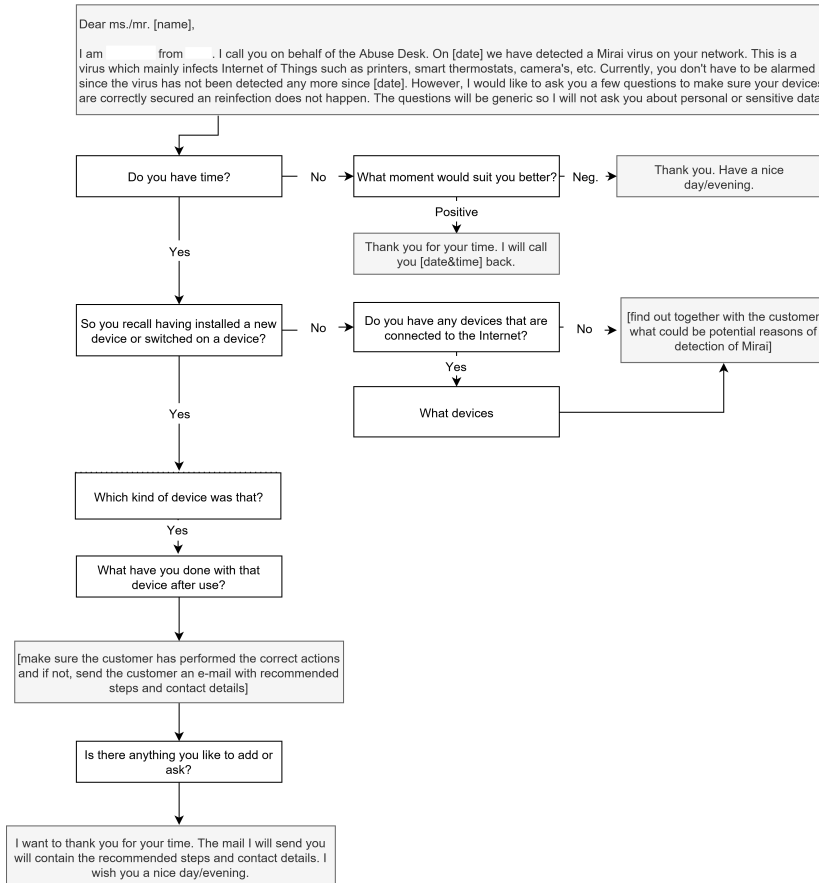


Figure 2: Survey protocol control group

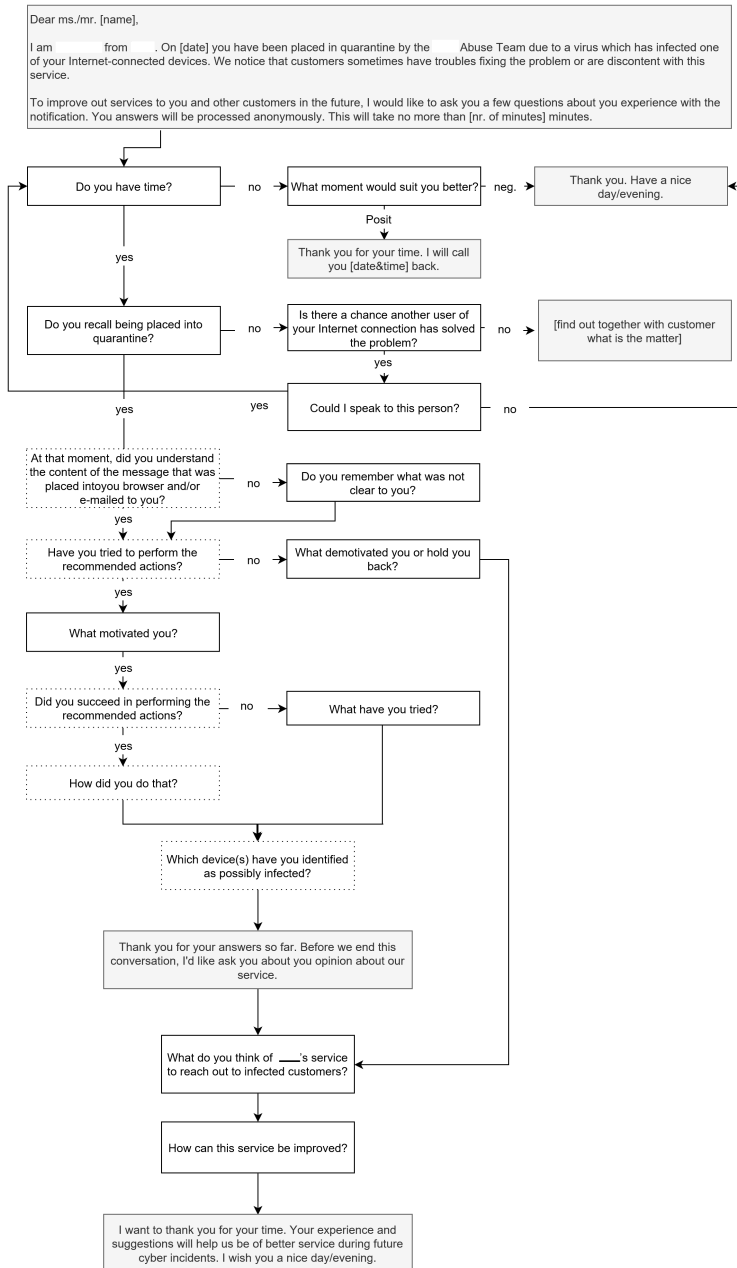


Figure 3: Survey protocol walled garden group

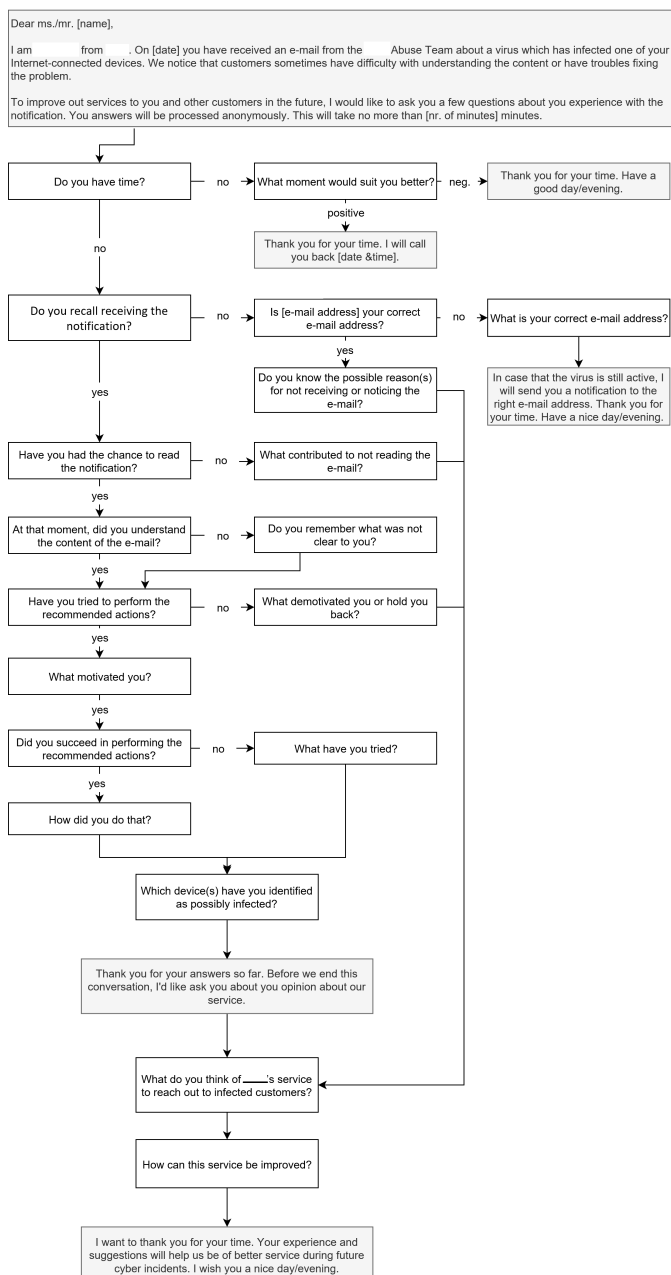


Figure 4: Survey protocol e-mail only group

A.4. NOTIFICATIONS

WALLED GARDEN

Illustration of walled-garden landing page displayed to consumers that were randomly assigned to the the walled-garden treatment group. The same content was also sent to consumers via email.

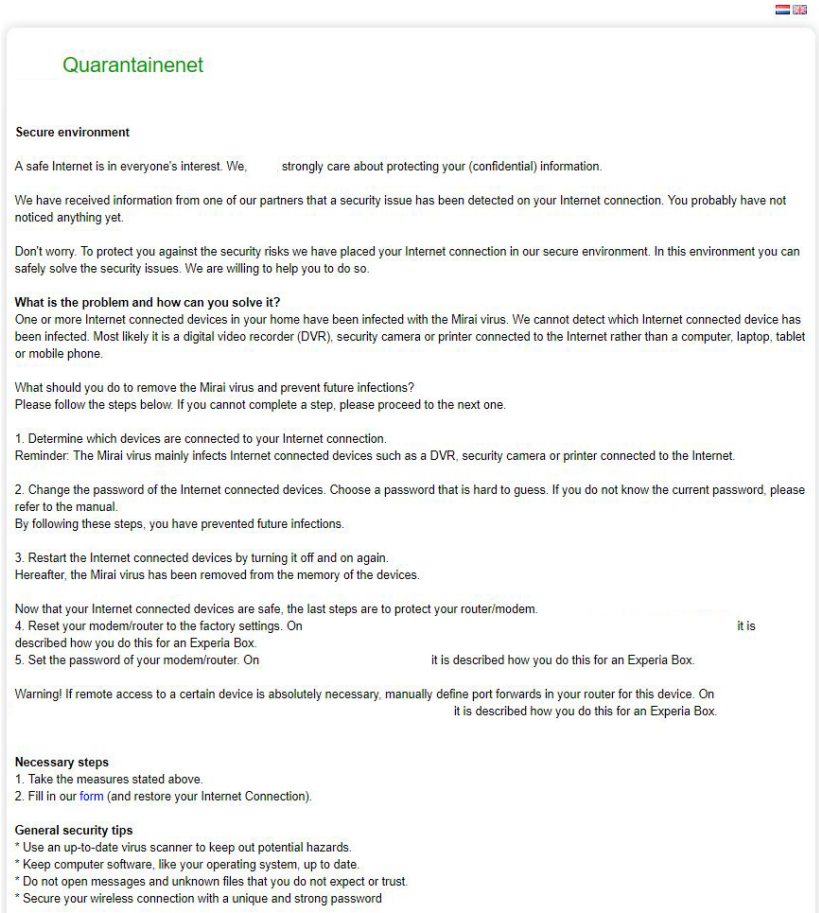


Figure 5: Landing page of walled garden

EMAIL-ONLY

Example of notification email sent to consumers randomly assigned to the email-only treatment group. The notification content essentially only differs with the previous example in that it omits statements about placing the recipient in a quarantine environment.

Dear Sir/Madam,

A safe internet is in everyone's interest. We, _____, strongly care about protecting your (confidential) information.

We have observed a security issue on your internet connection. You probably have not noticed anything, because it's about processes that run in the background.

One or more Internet connected devices in your home have been infected with the Mirai virus. We cannot detect which Internet connected device has been infected. Most likely it is a digital video recorder (DVR), security camera or printer connected to the Internet rather than a computer, laptop, tablet or mobile phone.

What should you do to remove the Mirai virus and prevent future infections? Please follow the steps below. If you cannot complete a step, please proceed to the next one.

1. Determine which devices are connected to your Internet connection. Reminder: The Mirai virus mainly infects Internet connected devices such as a DVR, security camera or printer connected to the Internet.
2. Change the password of the Internet connected devices. Choose a password that is hard to guess. If you do not know the current password, please refer to the manual. By following these steps, you have prevented future infections.
3. Restart the Internet connected devices by turning it off and on again. Hereafter, the Mirai virus has been removed from the memory of the devices. Now that your Internet connected devices are safe, the last steps are to protect your router/modem.
4. Reset your modem/router to the factory settings. On _____ it is described how you do this for an Experia Box.
5. Set the password of your modem/router. On _____ it is described how you do this for an Experia Box.

Warning! If remote access to a certain device is absolutely necessary, manually define port forwards in your router for this device. On _____ it is described how you do this for an Experia Box.

We ask you to take above steps within a day and to respond to this message. You can also ask additional questions in a reply to this email.

Kind regards,

_____ Abuse Team

Abuse team email

The _____ Abuse department deals with security incidents for _____. You can find more information about the Abuse department on: _____

Figure 6: Notification email sent to consumers in Email-only treatment group.

APPENDIX B

B.1. NOTIFICATION MESSAGE AND INSTRUCTIONS

Dear Sir/Madam [name],

We have discovered a security issue on your internet connection. We would like to resolve this issue together with you. The following sections explain how.

What is going on?

We have noticed that one or more internet-connected devices in your home have been infected with the mirai virus. While we cannot exactly detect which one of your connected devices has been infected, it is most likely a device such as a digital video recorder (DVR), security camera or printer connected to the Internet. Devices infected with the Mirai virus are typically **not** computers, laptops, tablets or mobile phones. The infection means that right now criminals have access to your infected device. This is putting you and other internet users at risk.

Tomorrow we will call you to resolve the issue

Our colleague, Mr. _____, will call you within a day to help you remove the virus. We gladly help you with this, as customers find it difficult to resolve the issue on their own. Moreover, the call will be a part of a scientific research that is executed together with _____ about the virus. This means we will ask you several questions to be able to help our customers better in the future.

Do you wish to remove the virus on your own?

Please let us know by a reply to this email or during the phone call. After that, please execute the following steps.

These are the steps needed to remove the virus

Step 1. Determine which devices are connected to your Internet connection. The Mirai virus mainly infects Internet connected devices such as a digital video recorder (DVR), security camera or printer connected to the Internet (not computers, laptops, tablets or mobile phones).

Step 2. Change the password of the Internet connected devices. Choose a password that is hard to guess. If you do not know the current password, please refer to the manual. By following these steps, you have prevented future infections.

Step 3. Restart the Internet connected devices by turning them off and on again. Hereafter, the Mirai virus has been removed from the memory of the devices.

Now that your Internet connected devices are safe, the last steps are to protect your router/modem.

Step 4. Reset your modem/router to the factory settings. On https://www._____.htm it is described how you can do this for an _____.

Step 5. Change the password of your modem/router. On https://www._____ it is described how you can do this for an _____.

NOTE: If remote access to a certain device is absolutely necessary, manually define port forwards in your router for the device. On https://_____/internet-9/port-forwarding-upnp-_____ it is described how you can do this for an _____.

Do you have any questions?

Please ask them in a reply to this email or during the phone call.

Kind regards,

Abuse Team
abuse@_____

The _____ Abuse department deals with security incidents for _____. You can find more information about the Abuse department on: https://www._____/abuse

Figure 7: Notification and opt-out invitation

B.2. THINK-ALOUD PROTOCOL

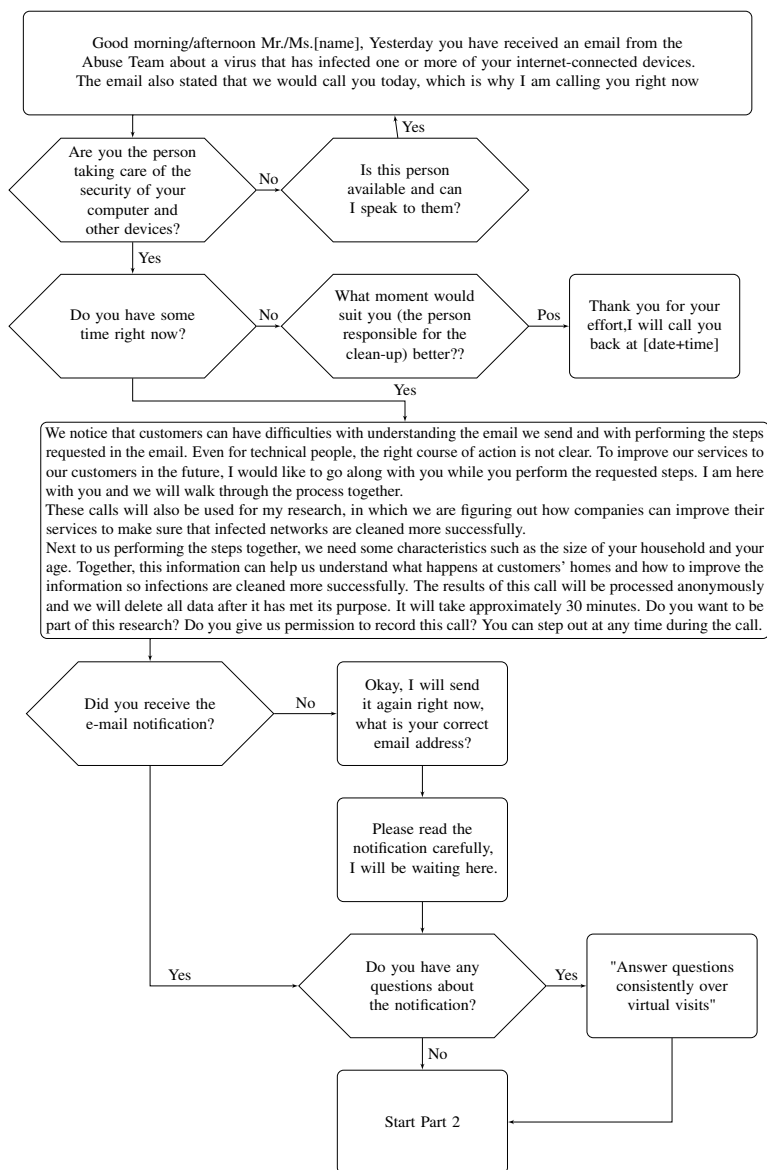


Figure 8: Think-aloud protocol - Part 1

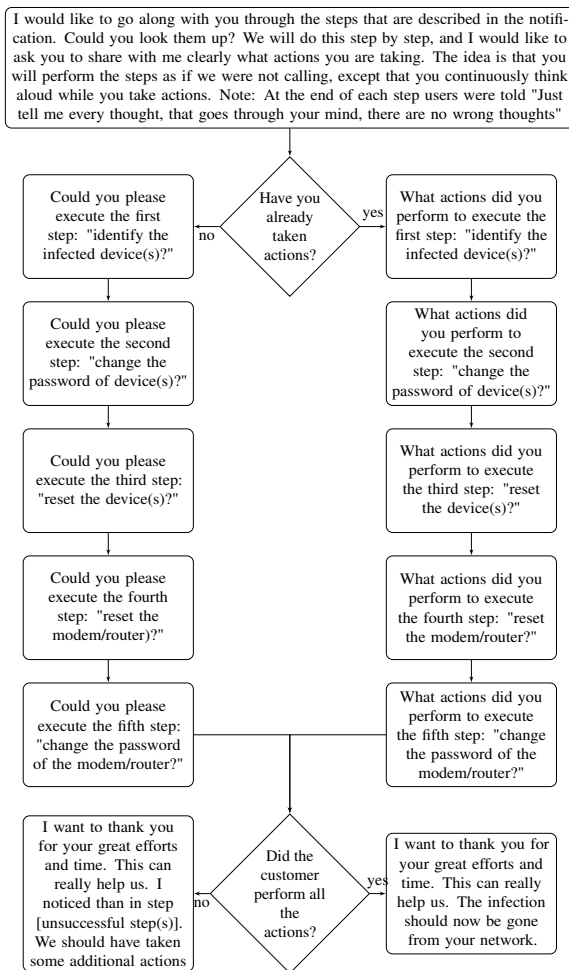


Figure 9: Think-aloud protocol - Part 2

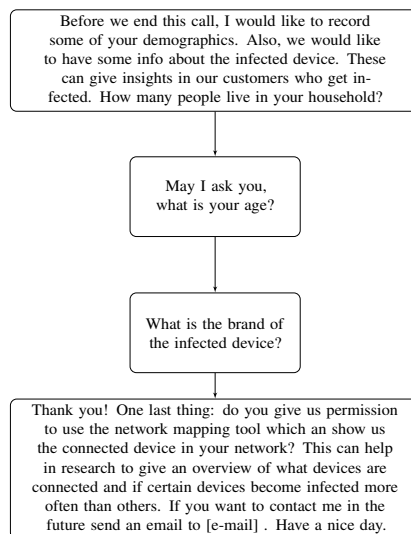


Figure 10: Think-aloud protocol - Part 3

APPENDIX C

C.1. EMAIL NOTIFICATION CONTENT

What's going on and how can I fix it?

A NAS from the supplier QNAP connected to your Internet connection is infected with the QSnatch malware. This infection poses a major risk to the safety of your files on the device. It is important to manually update your NAS operating system and malware remover app. Use the steps below:

Operating system:

- Go to the website: qnap.com/en-en/download
- Under "1 - Product type", select the option "NAS / Expansion" and select the number of slots present on the right.
- Under "3 - Model", select the type of NAS you are using.
- Under the "Operating System" tab, select the most recent version and download it via the "[REGION]" button.
- Open the NAS on your PC or Mac and choose firmware update, and then Manual update.
- Browse to the downloaded file and update the firmware / operating system.

Malware Remover app:

- Go to APP Center and choose "Malware Remover" and download it on your PC or Mac.
- Click on "manual update" in App center, browse to the download file and update the Malware Remover.
- Run a scan with the Malware remover.

What happens if I don't do anything?

The security problem on your Internet connection is a major threat. If you do not perform the steps or do not perform them correctly, we may place your Internet connection in our secure environment (quarantine). You can then temporarily make limited use of your Internet connection. By doing this we also protect your personal files and data.

Do you have any questions? Then you can ask this in a reply to this e-mail.

C.2. INTERVIEW QUESTIONS

	Did perform the steps	Did not/partially perform the steps
Check questions	Are you the person who manage the QNAP device?	Are you the person who manage the QNAP device?
	Do you use the device for business or private purposes?	Do you use the device for business or private purposes?
Opportunity	Did you receive the notification email?	Did you receive the notification email?
	Did you do the steps in notification email?	Did you do the steps in notification email? (if not) what did you do?
	Did you use any tools to perform the steps?	Did you lacked any tools to perform all of the steps?
	Did you have enough time to perform the steps?	Did you not have enough time to perform the steps?
	Was the location of the device or any of the tools you used an issue to access it?	Was the location of the device or any of the tools you used an issue to access it?
	Have any people helped you perform the steps?	Did any people try to help you perform the steps?
	Do some people you know have a strong opinion on performing the steps?	Do some people you know have a strong opinion on performing the steps?
Capability	Did you understand the steps?	Did you understand the steps?
	Did you find the steps challenging?	Did you find the steps challenging?
	Did you have any physical or bodily limitations that made the steps challenging?	Did you have any physical or bodily limitations that prevented you from finishing the steps?
	Can you give a rough indication of how much time it took to complete the steps?	Can you give a rough indication of how much time it took to complete what you did?
	Did you know what malware is?	Did you know what malware is?
	Did you know the difference between persistent and non-persistent malware?	Did you know the difference between persistent and non-persistent malware?
	Did you think you could perform the steps?	Did you think you could perform the steps?
	Did you have previous experience with IT systems?	Did you have previous experience with IT systems?
Motivation	Did you find the steps useful?	Did you find the steps useful?
	What do you think would happen if someone does not follow the steps?	What do you think would happen if someone does not follow the steps?
	Did you think you are responsible for performing the steps?	Did you think you are responsible for performing the steps?
	What did you feel while you performed the steps?	What did you feel when you received the notification email?
Exit question	Did an impulse helped you perform the steps?	Did an impulse prevent you from performing the steps?
	Is there anything that you would like to add that is relevant and we did not ask?	Is there anything that you would like to add that is relevant and we did not ask?

Note: Before the interview started, the researcher carrying out the interview took time to introduce himself, provided a description of the research, and asked for consent to proceed with the interview and data collection. Before the exit question, some demographic questions were asked, specifically self-reported gender and age.

APPENDIX D

D.1. MANUFACTURERS OFFERING SOFTWARE/FIRMWARE AND SECURITY ADVICE

	Manufacturer	Device	FW/SW	Password changing procedure	Advice to protect the device
1	ABUS	DVR	Yes	Yes	No
2	Advanced Multimedia Internet Technology (AMIT)	WIP-300 Router	No	Yes	No
3	ASUS	RT-AC5300	Yes	Yes	No
		RT-N10U	Yes	No	No
		RT-AC58U	Yes	Yes	Yes
		RT-N10 + B1	Yes	No	No
		RT-AC54U	Yes	Yes	Yes
		RT-AC87U	Yes	Yes	Yes
		RT-N14U	Yes	Yes	No
		RT-N13U.B1	Yes	No	No
		RT-G32	Yes	No	No
		RT-N10	Yes	No	No
		DSL-N10	Yes	No	No
		WIRELESS-AC1200	Yes	Yes	Yes
4	AVM GmbH	FritzBox Router	Yes	Yes	Yes
5	AirTies	Air4920-2 SetTopBox	No	No	No
		Air7120 SetTopBox	No	No	No
6	Amlogic	SetTopBox S905L	No	No	No
7	Asustor	NAS	Yes	No	No
8	Avtech	IP Camera, DVR	Yes	Yes	No
		IP Camera	Yes	Yes	Yes
9	Bab Technologie	Unknown	NA	NA	NA
10	Beijer Electronics	QTERM Panel	No	No	No
11	Broadcom	BCM Router	No	No	No
12	Ceru Co. Ltd	vu+ Solo2	No	No	No
13	Cisco	Docsis Gateway	No	No	No
14	D-Link	Router	Yes	Yes	Yes
15	Devolto	Microlink Dlan Wireless	Yes	No	No
16	Digicom	RAW300L-A05 Router	Yes	Yes	Yes
17	Domoticz	Home Automation	Yes	Yes	No
		Domoticz Machinon	Yes	Yes	No
18	DrayTek	Vigor 2860 Router	Yes	Yes	Yes
		Vigor 2925 Router	Yes	Yes	Yes
		Vigor 2760 Router	Yes	Yes	Yes
		Vigor 2960 Router	Yes	Yes	Yes

		Vigor 2926 Router	Yes	Yes	Yes
		Vigor 2133F Router	Yes	Yes	Yes
		Vigor 2862 Router	Yes	Yes	Yes
19	Dream Multimedia	Dreambox DVB Satellite	No	No	No
20	Fibaro	Home Centre	Yes	No	No
21	Flying Voice Technology	FWR9601 VoIP Router	Yes	No	No
22	Foscam	Foscam	Yes	Yes	Yes
23	Freebox	SetTopBox	Yes	No	No
24	GNSS	Receiver Net-G5 GNSS	Yes	No	No
25	Grandstream	UCM6202 IP PBX	Yes	Yes	Yes
26	Hichan Technology	Router WiDisk	No	No	No
27	HikVision	IP Camera	Yes	Yes	Yes
		DVR	Yes	Yes	Yes
28	Hisilicon	Hi3798MV300 SetTopBox	No	No	No
29	Huawei	Router	Yes	Yes	No
		SetTopBox	No	No	No
		Home Gateway	No	Yes	Yes
		HG659	No	Yes	Yes
30	Inim Electronics	Smartlan Fire Control System	No	No	No
31	Innbox	VDSL2 modem	No	No	No
32	Interlogix	TruVision NVR	Yes	Yes	Yes
33	Level One	WBR-6005 Router	No	Yes	Yes
34	Lifetrons	FG1060N Wifi Router	No	No	No
35	Linksys	Router	Yes	Yes	Yes
		Linksys LRT214	Yes	Yes	Yes
36	MAGINON	Camera, camcorders, other electronics	Yes	Yes	Yes
		IPC-250HDC	Yes	Yes	Yes
		Security Camera	Yes	Yes	No
37	Merit Lilin	NVR	No	Yes	No
38	MikroTik	Router	Yes	Yes	Yes
		Router v6.12	Yes	Yes	Yes
		Router v6.43.12	Yes	Yes	Yes
39	Netcomm	VDSL2 N300 WiFi Router	Yes	Yes	No
40	Netis	Router	Yes	No	No
41	Opendreambox	SetTopBox	No	No	No
42	Phicomm	Router	No	Yes	Yes
43	QNAP	QNAP QTS	Yes	Yes	Yes
		Network Attached Storage	Yes	Yes	Yes
		QNAP QTS 4.3.3.1098	Yes	Yes	Yes
		QNAP QTS 4.4.2.12.62	Yes	Yes	Yes
		QNAP QTS 4.3.4.1129	Yes	Yes	Yes
		QNAP QTS 4.2.6	Yes	Yes	Yes
		QNAP QTS 4.2	Yes	Yes	Yes
44	Reolink	NVR	Yes	Yes	No
45	Ricoh	Aficio MP 301 Printer	No	No	No
46	Samsung	DVR	Yes	Yes	No
47	Sansco	NVR Security Camera	No	Yes	No
48	Siera	Siera Panther DVR	No	No	No
49	Sompy	Alarm System	No	No	No
50	Sony	Ipela SNC-CH160	Yes	Yes	Yes
51	STMicroelectronics	Unknown	NA	NA	NA
52	Strong	Extender 1600	Yes	No	No
53	Synology	Disk Station	Yes	Yes	Yes

		Disk Station DS916	Yes	Yes	Yes
54	TOTOLink	Router	Yes	Yes	No
55	TP-Link	Router	Yes	No	Yes
56	Tecom	AH2322 ADSL Router	No	No	No
57	Ubiquiti	Aircube AC	Yes	No	No
58	Uniview	Unv IP Camera	No	No	No
59	Upvel	UR 313N4G Router	Yes	Yes	No
		UR-321BN Router	Yes	Yes	No
60	VACRON	NVR	Yes	No	No
61	Vimar	Elvox Video Door entry	Yes	No	No
62	X10 Wireless Technology Inc	IP Camera AirSight Xx34A	No	No	No
63	XPO Tech	ZEM560 Fingerprint	No	No	No
64	Xiong Mai	White labeling DVR, White labeling	Yes	No	No
		NVR			
		DVR	Yes	No	No
		NAS	No	No	No
65	ZKTeco	ZEM560 Fingerprint	No	No	No
		ZMM220	No	No	No
66	ZNDS	Smart TV Box	No	No	No
67	ZTE	Router	No	No	No
		F620V2 Router	No	No	No
68	Zhejiang Dahua Technology Co., Ltd.	IP Camera (IR PTZ Dome Camera)	Yes	Yes	Yes
		IP Camera	Yes	Yes	Yes
69	Zhone Technologies	ZNID-GPON-2426A-NA Router	No	Yes	No
70	Zyxel	ADSL gateway	No	No	Yes
		WAP5705 Media Streaming Box	No	Yes	Yes
		NSA325 v2	Yes	Yes	Yes

APPENDIX E

E.1. PUBLIC DNS RESOLVERS CLASSIFICATION

Table 3: Public DNS resolvers classification

Classification	Public DNS resolver name	IP addresses
Protective DNS (PDNS)	114 DNS	114.114.115.115
	AliDNS	223.6.6.6
	Alternate DNS	198.101.242.72
	Baidu DNS	180.76.76.76
	CleanBrowsing	185.228.168.9
	Comodo Secure DNS	8.26.56.26
	DNS PAI	101.226.4.6
	DNSPod	119.29.29.29
	Green Team DNS	81.218.119.11
	Neustar	156.154.70.1
	One DNS	117.50.10.10
	OpenDNS	208.67.222.222
	SafeDNS	195.46.39.39
Possible Protective DNS	Cloudflare	1.1.1.1, No Malware: 1.1.1.2 No Malware and adult content: 1.1.1.3
	Yandex	Basic: 77.88.8.8,77.88.8.1, Safe: 77.88.8.88,77.88.8.2, Family: 77.88.8.7,77.88.8.3
	Quad9	9.9.9.9 No filtering: 9.9.9.10
No Protective DNS	CNNIC SDNS	1.2.4.8
	DNS.Watch	84.200.69.80, 84.200.70.40
	Freenom World	80.80.80.80
	Google Public DNS	8.8.8.8
	Hurricane Electric DNS	74.82.42.42
	Open NIC	96.90.175.167
	Oracle Dyn	216.146.35.35, 216.146.36.36
	Quad101	101.101.101.101
	Uncensored DNS	91.239.100.100
	Verisign OpenDNS	64.6.65.6
No information	Free DNS	45.33.97.5
	Level 3	209.244.0.3
	puntCAT	109.69.8.51

E.2. FOCUS GROUPS AND PILOT

Before launching the Prolific survey, we performed two focus groups. The first included five participants from our computer science department and the second had four people without a background in computer science. Thanks to the first focus group, we reduced the survey size and switched from a conjoint analysis to a standard survey because participants said it was easy to flick through the options. We toned down technical explanations of PDNS and further explanations were added to the questions after the second focus group. We ran a pilot with 10 participants in Prolific to check everything was working fine. We didn't change any questions, thus we used pilot data in the study's results.

E.3. DNS MEASUREMENT

A DNS measurement similar to the APNIC data collection was integrated into the Prolific survey. We included a Javascript that was triggered when participants submitted their unique Prolific ID. The Javascript fetched 'https://prolific ID + .[DOMAIN NAME UNDER OUR CONTROL]'. We recorded their resolver's IP addresses to determine if they were using PDNS or not. We mapped participants' IP and their resolvers' IP to ASes using Pyasn [27]. Out of the 295 participants, we obtained DNS logs for 285 of them.

E.4. SURVEY INSTRUMENT

<https://doi.org/10.4121/22232911.v1>

E.5. INTERVIEW PROTOCOL ISP CASE STUDY

Informed consent

- 1) What kind of Internet-connected devices do you own?
 - 2) Do you think that your online devices are secure against being abused? Why or why not?
 - 3) What do you think can be the consequences of abuse of Internet-connected devices?
 - 4) Who do you feel should be responsible for the security of Internet-connected devices?
 - 5) Do you use any security software or services or other security precautions to protect your Internet-connected devices?
 - **If 'Yes' answered to question 5:** 6) What kind of security measures do you use?
 - **If 'No' answered to question 5:** 6) Why you do not use any security measures?
 - 7) Did you enable the [ISP name] [service name]?
 - **If 'Yes' answered to question 7:** 8) Why did you enable the [service name]?
 - **If 'No' answered to question 7:** 8) Why you did not enable the [service name]?
 - 9) Do you think there could be any drawbacks associated with the use of services like [service name]?
 - 10) How do you feel about your ISP offering the [service name]?
- Demographics questions

E.6. ENTERPRISE INTERVIEWS

Informed consent

- 1) What is your role in this organization?
- 2) What is your organizations' core business?
- 3) How many employees does your organization have?
- 4) What network security concerns does your organization have?
- 5) Do you have network security policies and measures that address the network security concerns that your organization has?
- 6) How does your users' activities relate to those policies and security measures?
- 7) What kind of DNS resolver does your organization use?
- 8) Does your organization use any form of filtering in the network at DNS level?
- 9) Are you aware of services that filter malicious domains?
- **If the organization uses Protective Domain Name System:**
 - 10) Why does your organization use these subsets of measures [mentioned in question 5] and PDNS?
 - 11) Why did your organization choose to use PDNS as an additional measure?
 - 12) How is PDNS used in your organization?
 - 13) How costly it is to use PDNS versus other security measures?
 - 14) Which results of the use of PDNS are most valuable? How often does this occur?
 - 15) Have your organization ever had any problems in the operation of the network due to the use of PDNS?
 - 16) What do you think about government initiatives about PDNS? (e.g. CIRA Canadian shield, The United Kingdom, Australia, and DNS4EU)?
 - 17) Will your organization change your current PDNS for one provided by the government?
- **If the organization does not use Protective Domain Name System:**

(Note: Definition of PDNS was provided in case the participant didn't know what PDNS was)

 - 10) Could your organization consider using something like PDNS?

- 11) Do you think that a service such as PDNS can be an addition to your security measures?
- 12) What factors will your organization consider to use a service such as PDNS as an additional measure?
- 13) How costly do you think the use of PDNS can be versus other security measures?
- 14) Which results of the use of a service such as PDNS could be most valuable to your organization?
- 15) Could you foresee any problems with the use of PDNS in the operation of the network?
- 16) What do you think about government initiatives about PDNS? (e.g. CIRA Canadian shield, The United Kingdom, Australia, and DNS4EU)
- 17) Will your organization consider using a PDNS service provided by the government?

E.7. EXPERTS INTERVIEWS

Informed consent

- 1) What is your experience with DNS?
- 2) What do you understand as Protective Domain Name System?
- 3) What do you think of Protective Domain Name System for security purposes?
- 4) How Protective Domain Name System is different from other current available security solutions?
- 5) Do you have any concerns about the operation of Protective Domain Name System ? Prompts: (i) Who should be offering this service? (ii) Who should be using this service? (iii) What factors should be considered in order to adopt PDNS? (iv) Pros and cons / factors for success and failure
- 6) What do you think about government initiatives about Protective Domain Name System? (e.g. CIRA Canadian shield, The United Kingdom, Australia, and DNS4EU)

E.8. QUALITATIVE CODING

Table 4: Summary of qualitative coding scheme ISP interviews

Themes	Code examples	Respondents n=24
Concerns about the service	Privacy, data usage, cost	21 (88%)
Reasons for not adoption	Other SW to block malware, cost once enabled	15 (62.5%)
Belives on abuse	Identity fraud, phishing, spread malware, data stolen	15 (62.5%)
Trust	Trust ISP, distrust email	11 (46%)
Reasons for adoption	Useful service, prevent malware, ISP advice	9 (37.5%)

Table 5: Summary of qualitative coding scheme enterprise interviews

Themes	Code examples	Respondents n=12
Awareness of PDNS	Knows about PDNS, does not know about PDNS	12 (100%)
Concerns about PDNS	Privacy, false positives	9 (75%)
Government PDNS	Welcome government initiatives, Do not welcome government initiatives, useful to have options	9 (75%)
Factors to consider for adoption	Layered security, threat model	7 (58%)
Reasons to implement PDNS	Global TI, reputation	6 (50%)

Table 6: Summary of qualitative coding scheme experts interviews

Themes	Code examples	Respondents n=9
Factors for adoption	Performance, users awareness, organizations' security strategy	9 (100%)
Provider	Who should offer PDNS, PDNS provider, gov as provider	9 (100%)
Limitations of PDNS	What can go wrong, complementary solution	9 (100%)
Types of blocking	Legal basis blocking, blocking for security purposes, benign content blocking	8 (89%)
PDNS vs other security measures	DNS path broken, no installation, all devices protected	8 (89%)
Privacy	Data sharing, data monetization, privacy	7 (78%)
Transparency	Who decides what to block, transparency	6 (67%)
Centralization	Options to choose, diversification	6 (67%)

E.9. VARIABLES INCLUDED IN THE FINAL ORDINAL REGRESSION MODEL

Reference category	Variables	Explanation of coding	Survey questions
	Concerns	Factor analysis	Q24,Q25,Q26
	Perceived vulnerability	Factor analysis	Q6,Q7,Q8
	Perceived severity	Factor analysis	Q9,Q10,Q11,Q12
	Useful	Continuous scale	Q19
No security installed by themselves	Install security themselves	True if participant recall setting up security features in his internet-connected devices by himself and did not provide any other answer	Q13
Do not use any security tool	Use other security measures	True if the participant uses Antivirus or Firewall or Ad blocker or any other tools to protect his devices and did not answer that he does not implement any security tool.	Q14
Does not use parental control	Use parental control	True if participant uses parental control	Q15
Not aware	Aware	True if participant heard before of a similar service like PDNS only	Q17
	Unsure	True if the participant was not sure of hearing of a similar service like PDNS only	
Not willing to pay	Willing to pay	True if participants were willing to pay for PDNS service.	Q34
Government provider	Commercial provider ISP provider Other provider	True if commercial provider was selected and not government or ISP provider or other provider. True if ISP was selected as provider and not government or commercial provider or other provider. True if other provider was selected and not government,or ISP provider or commercial provider.	Q30
Control variables:			
Reference category	Variables	Explanation of coding	Survey questions
Male	Female Other genders	True if participant identify as female and not as male or other genders True if participant identify as other gender and not as Male or Female	Q36

E.10. TOP 20 COUNTRIES WITH PDNS USAGE

Table 7: Top 20 countries with the highest percentage of DNS queries answered by PDNS (Period: January to June 2022)

cc	avg daily queries	Internet users	% sampled	Non Public DNS resolvers			Public DNS resolvers			
				% Same AS	% In country	% Out country	% PDNS	% No PDNS	% Possible PDNS	% No Info
IL	53,235	7002759	0.76%	50.4%	2%	0.3%	34%	12.1%	1.1%	0.1%
AF	10,963	9327489	0.12%	25%	0%	5%	25%	40%	4%	1%
CY	15,205	1011831	1.5%	53.4%	12.4%	0.4%	23%	8.5%	1.8%	0.5%
ME	18,283	449989	4.06%	54.4%	0.1%	1%	22.3%	22%	0.2%	0%
TZ	45,146	23142960	0.2%	44%	10%	0%	9%	36%	1%	0%
GE	38,407	32543600	0.118%	49%	25%	0%	8%	13%	5%	0%
ZM	22,695	9870427	0.23%	3%	0%	0.3%	7.3%	72.4%	0%	17%
NG	213,900	126078999	0.17%	65%	3%	2%	7%	23%	0%	0%
IR	94,366	67602731	0.14%	26%	9%	31%	3%	8%	22%	1%
AL	57,115	2160000	2.64%	67.1%	0.4%	0.1%	2.6%	25.3%	4.3%	0.2%
US	871,976	313322688	0.28%	62.4%	10%	2.1%	2%	20.1%	3%	0.4%
EG	475,809	49231493	0.97%	68%	13%	0.5%	2%	16%	0.5%	0%
VN	152,645	84883000	0.18%	68.4%	0.8%	0.2%	1.1%	27.5%	1.9%	0.1%
ID	1,320,259	212354070	0.62%	68%	17%	0%	1%	12.4%	1.6%	0%
BR	525,834	150457635	0.35%	48.3%	15%	1.4%	1%	28.3%	6%	0%
TR	228,978	69107183	0.33%	53%	31%	0%	1%	14%	1%	0%
UA	158,809	40912381	0.39%	68%	4%	1%	1%	20%	6%	0%
PH	595,824	95200000	0.63%	45%	34%	0.2%	0.4%	18%	2.3%	0.1%
IN	3,207,855	755820000	0.42%	58.1%	27.2%	0.1%	0.2%	14%	0.3%	0.1%
BD	790,017	117310000	0.67%	58%	5%	0.4%	0.2%	32.4%	4%	0%

Note: % **Same AS**: Percentage of average daily queries which resolvers ARE in the same AS as the users and NOT known public DNS resolvers. % **In country**: Percentage of average daily queries which resolvers are NOT known public DNS resolvers and NOT in the same AS as the users, but ARE geolocated in the same country as the user.

% **Out country**: Percentage of average daily queries in which resolvers are NOT known public DNS resolvers and NOT in the same AS as the users but, and NOT geolocated in the same country as the user. % **Public DNS resolvers**: percentage of average daily queries which are answered by resolvers as categorized in Table 3.

E.11. ORDINAL LOGISTIC REGRESSION

Figure 11 displays on the left side the variables with beta (β) which is the estimated regression coefficients of the variables (all the beta values in the graph are significant at $p < 0.1$), and their standard error. On the right side the odds ratio (OR), which is the exponentiated regression coefficient, and their confidence interval.




Figure 11: Significant predictor variables

ACKNOWLEDGEMENTS

I am deeply grateful when I think of all the people who have supported me along my PhD journey. Many people, in one way or another, have given me one of the most precious gifts I could ever receive, their time. I realize that one section could never do justice, but I will make an attempt.

First, I would like to express my gratitude to Michel van Eeten, my promotor. I appreciate your commitment to our biweekly meetings despite your busy agenda. I have always been impressed by your dedication and passion for your work, and I have learned a great deal about storytelling, consistency, and formulating sharp answers thanks to you. Although I am still a work in progress. I will always treasure the warmth and hospitality you and your family showed me over Christmas when I was alone in The Netherlands and for making me feel at home.

My daily supervisor and promotor, Carlos, also deserves my heartfelt thank you. I appreciate your support, which dates back to the first day of my PhD journey. I appreciate you taking the time to answer my questions and for your insights about the playfield of scientific research. Your different perspectives and assessment of our work were always insightful and predicted potential problems before they arose. I value the freedom you gave me to learn from my experiences and mistakes. I am grateful for the endless Slack chats, your help and encouragement, and the trust you placed in me.

I would like to thank Raymond Teunissen, Michel Zoetebier, and Dennis van Beusekom, who were all instrumental in making the ISP case studies possible. Thank you for serving as external project supervisors and handling the administrative details, which allowed us to launch these research projects.

Also, my gratitude goes out to Susanne Verstegen, Swaathi Vetrivel, Brennen Bouwmeester, Max Fukkink, Ralph van Gorp, and Kevin Siu, all of whom were master's students with whom I collaborated. In addition to being a great learning experience, being involved in your master's thesis projects was one of the highlights of my PhD experience.

Many thanks to all members of the Economics of Cybersecurity group. Radu, Veerle, Evi, Aksel, Xander, Fieke, Mannat, Nishant, Donald, Yury, Savvas, Simon, Seda, Rolf, Arwa, Mathew, Sandra, Swaathi, Lorenz, and others who were formerly a part of the group: Boy, Arman, Natalia, Qasim, Bernardus, Ugur, Tobias, Wolter, Kate, and Maaïke. I love your positive attitude, great questions, and the stimulating conversations we have had through the years. To new group members, Szu-Chun, Hugo, Kelvin, Cécile, Ron, Gerbrand, although we had not much chance to share, I can tell you that you are joining a great team, and I wish you all the best in your research. Special thanks to Gerbrand for helping me as a second reviewer of the Dutch translation of the executive summary of this dissertation.

A heartfelt 'thank you' to Arman. Thank you for your kind words and helpful advice as I began my PhD journey. You were instrumental as the senior PhD I met when I started my journey.

I would like to thank Simon for all the helpful discussions and guidance. Thanks for all the effort you put into our work despite your hectic schedule. Thanks for your unwavering commitment. You are

a genuine academic, as I am fond of reminding you.

I want to express my gratitude to Yury as well. Thank you for our coffee breaks and the insightful conversations about what is good research. Also, thank you, Savvas, for introducing me to memes; from now on, there is no going back in presentations. I shall always use a meme. Natalia, thank you for being vital to the group's enjoyable social activities. Rolf, thank you for all of your guidance on formalities for my defense, organizing my mock defense, and your advice on potential job paths after my PhD.

My deep appreciation to Sandra & Tadashi, Lorenz & Ketii, Swaathi & Omkar. Thanks for the time you and your loved ones have taken to get to know me outside of office hours and offer me your friendship.

I have found close-knit friends at TPM, Matt, Boy, Maria José, and Arwa, I give thanks to the corona times for bringing us together.

Boy and Matt, I appreciate you being there for me on some of the most challenging days of my PhD, encouraging me and listening to me without passing judgment. We may have differing opinions on many topics, but we can all agree that you are such good friends.

Boy, I appreciate our walks around the lake, your unique perspective on life, and the courage you display in advocating for your beliefs.

Thank you, Matt, for always answering my 'cat eyes' requests. Thanks for always taking time from your busy day to help me with things like my executive summary translation to Dutch, debugging code, and so on. I appreciate the time we spent chatting before starting to work, I will miss that.

María José, thank you for being there for me and being such an amazing friend, and for all the wonderful dinners, walks around the lake, and chats we have shared. Your search for yourself and growth has always inspired me.

Arwa, thanks for your friendship, kindness, and support. You are sweet, thoughtful, brave, and the best host. I am lucky to have a princess friend (I couldn't stop myself). Thank you for adding Yoda as a new source of entertainment to our lives.

I want to express my appreciation to the entire staff of the Organization and Governance department. Special thanks to the secretaries, Joy, Wendela, and Jolanda, who worked tirelessly behind the scenes to make my administrative tasks easier.

Joyce van Velzen, I appreciate your understanding as I worked to fill my TIM hours, and I cherish the plant you gave me as my first personal gardening endeavor in the office.

Patricia Carrion and Ismail Yetin, I'll never forget your kindness and guidance upon my arrival in the Netherlands in 2016. Your assistance in beginning my master's degree set me on this path.

To my master's thesis advisor, Hadi Asghari, I like to say thanks. You were instrumental in my decision to continue my education and earn a PhD. Under your guidance, I grew to like conducting research.

I would like to thank my ISC family. Ruben and Henk thank you for sharing your wisdom and experience about life, for checking on me when I was sick, and for all your messages of support through these years.

Pavel & Claudia, Kate and Nikolas, Gijs & Claudia, Sofia, Elías, you have a very special place in my heart. Thank you for letting me in your home and hearts, and for all the support, hospitality, and kindness. I appreciate you sharing your family's love with me.

Francesco & Mariana, Pietro, Alessandro. I remember perfectly the warmth I felt since the first time I went to your home. Thank you for your friendship during all this time.

Ella & Jose, Pablo, Isabel, Mikha, Ernes & Pablo, Mona, Consta & Mithun, Matilde. Thank you all for welcoming me into your lives and allowing me to see you grow as families and be part of your journey.

Nicola & Roberta, Mía, Lea, Santiago, Maria Parra, Mohan, Fr Avin, Rev Taco, Rev Waltraut, Theo, Jane, Wisdom, William, Ashish & Hannah, Arun & Merin, Maia & Italo, James, Carmen & Luca, Monica. Thank you for always making me feel at home, for the coffee times, potlucks, and for your friendship. Jane I appreciate our garden dates, you are a wonderful kind soul. All the ones that are not in The Netherlands anymore: Praveen, Shyrle, Silvia, Amit, Sarita & Rogelio, Ruby & Eduardo, Mathew A., Priya, James Z. Forgive me if I forget some names. Thanks all for being my home away from home.

I would like to thank the mandatory group. Since we first met in our master's program, I want to say how much I appreciate the many enjoyable dinners and time we have spent together. You all are part of the family I have in The Netherlands now.

Astrid, Zarife, Lianne, Bram, and Milan. It is not easy to make a Dutch friend, but once you do, you have got one for life. You are definitely one of those. A special thank you to Astrid for the many dinners we have shared together to catch up on each other's lives and our growth.

Manuela, Agnelo, Vimal, Mihir, we have been in closer contact over the years due to our prolonged residence in Delft. I feel fortunate to have witnessed your development alongside mine.

Although I don't see you very often, Rahul, Tanya, Akhil, I will always hold a very special place in my heart for you. To this day, I will never forget the time we spent together studying for Technology Dynamics, you are part of where this journey started. Rahul, thank you for allowing me to share joyous occasions of your life such as your wedding with Eva, and Dhyan's birthdays.

Elisa, Renita, Ifa, Lina, Andres, Efrain. Thanks all for being part of this journey. I know that friends like you are hard to come by.

Alma Iris, Ever, Violeta, Ly, Ronald, Mey, Ines, Maria del C, Darling, Lux, Indira, and Raulito, all of whom I've known for years. All of you have been wonderful friends and supporters throughout my life. No matter how infrequently we communicate, know that you are never far from my thoughts.

I would like to express my gratitude to Daria and Ana I, two remarkable women I have met along this way. Thank you for helping me become a better version of myself.

To all my extended family, Miguel Angel Rodríguez Zeledón, Cristina & Valentin, Miguel Angel Valentin, Aaron, Cris, Fer, Leslie, Marco Emilio, Camilo, Miguel & Marbelí, Julio Miguel, Angelita, Christopher, Angela M. & Jason, Darling M. & Paul, Edwin M., Marco Antonio, Freddy M., Yadira J., Yeni & Medardo, Leonor, Victoria Morales, Victoria Rivera, Victoria Rodríguez, Lucy González, Darling Centeno. Thank you all for your unwavering support during this journey. I am grateful to have you all in my life.

To Valen, I must express my deepest gratitude. To me, you are just like a brother. Words cannot express how much your support meant to me.

To my mom, Evelyn, thank you for always believing in me and allowing me to pursue my goals and dreams. I could not have done this without you! You have taught me a great deal about perseverance and grit. This accomplishment is dedicated to you. The words I write here cannot express all my gratitude and love for you.

A mi madre, Evelyn, gracias por creer siempre en mí y permitirme perseguir mis metas y sueños. ¡No podría haber hecho esto sin tí! Me has enseñado mucho sobre la perseverancia y la determinación. Este logro esta dedicato a tí. Las palabras que escribo aqui no pueden expresar toda la gratitud y el amor que siento por tí.

To my 'little' sister Linda, I am extremely grateful to have a sister like you. You are a determined, brave, smart, sensitive, and strong woman. To me, there are no limits to what you can accomplish in life, but remember that you are already enough! I will always be there to support you no matter what.

Last but not least, I want to express my gratitude to Bayardo. Thank you for being there for me throughout the years, and for your love, friendship, support, understanding, and encouragement. It has been a long journey, but now we are both doctors.

Elsa Rebeca Turcios Rodríguez
Nicaragua, December 2022

AUTHORSHIP CONTRIBUTIONS

The dissertation is based on five peer-reviewed papers resulting from collaboration with multiple co-authors. In each of these studies, I was fortunate to receive valuable feedback and a variety of contributions from my co-authors. In the next paragraphs, I will describe each of their individual contributions to each study.

For the first study (see Chapter 2), my co-authors, Arman Noroozian, Carlos Gañán and Michel van Eeten, have helped with improving the draft, refining its arguments, proofreading and polishing the text. Daisuke Inoue and Takahiro Kasama provided access to the network telescope data and did the data processing of it. Susanne Verstegen carried out the survey. Analysis of the underlying data and statistical modeling used for this study was performed by me.

For the second study (see Chapter 3), Brennen Bouwmeester carried out the think-aloud protocol and data collection and did the initial coding of the thematic analysis. Michel van Eeten helped with the results section as the second coder for the thematic analysis. Simon Parkin sharpened the paper's argumentation and discussion. Carlos Gañán aided in revising the document, checking for errors, and formatting the text and tables. I was one of Brennen Bouwmeester's advisors during the course of this research (this publication is the product of his master's thesis), and I was involved in the research design.

For the third study (see Chapter 4), my co-authors Simon Parkin, Michel van Eeten and Carlos Gañán have all helped with improving my drafts and sharpening the paper's structure. The discussion was refined with help from Simon Parkin. The interviews were carried out by Max Fukkink and he and I independently coded the interviews.

For the fourth study (see Chapter 5), my co-authors Carlos Gañán and Michel van Eeten, helped to improve the draft, refine its arguments, proofread, and polish the text. Arman Noroozian helped me to optimize the script for the data collection and focused on improving the text of the methodology section. The manual labeling of the ground truth to identify manufacturers, data analysis, and checking manuals and websites of manufacturers were performed by me.

Finally, the fifth study (see Chapter 6), builds on a greatly valued discussion with my co-authors, Radu Anghel, Simon Parkin, Michel van Eeten and Carlos Gañán and their help sharpening the paper's structure and arguments. Also, Radu Anghel did the configuration of the DNS server for the measurement incorporated in the survey. The data collection and analysis were handled by me.

I owe a great deal to the encouragement and guidance of my promotor Michel van Eeten and co-promotor Carlos Gañán who patiently guided me through the process of all these studies. Simon Parkin, you are greatly appreciated for your ability to enhance the articles' discussion sections even before they are written. To my other co-authors, I am grateful for your support, ideas, and feedback.

LIST OF PUBLICATIONS

- **Rodríguez, E.**, Verstegen, S., Noroozian, A., Inoue, D., Kasama, T., van Eeten, M., & Gañán, C. H. (2021). “User compliance and remediation success after IoT malware notifications”. In *Journal of Cybersecurity*, 7(1), tyab015.
- **Rodríguez, E.**, Noroozian, A., van Eeten, M., & Gañán, C. (2021). “Super-Spreaders: Quantifying the Role of IoT Manufacturers in Device Infections”. In *20th Annual Workshop on the Economics of Information Security (WEIS 2021)*.
- Bouwmeester, B., **Rodríguez, E.**, Gañán, C., van Eeten, M., & Parkin, S. (2021). “The Thing Doesn’t Have a Nam”: Learning from Emergent Real-World Interventions in Smart Home Security. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)* (pp. 493-512).
- **E. Rodríguez**, M. Fukkink, S. Parkin, M. van Eeten & C. Gañán, “Difficult for Thee, But Not for Me: Measuring the Difficulty and User Experience of Remediating Persistent IoT Malware” *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*, 2022, pp. 392-409, doi: 10.1109/EuroSP53844.2022.00032.
- **E. Rodríguez**, R. Anghel, S. Parkin, M. van Eeten & C. Gañán, “Two Sides of the Shield: Understanding Protective DNS adoption factors” (*USENIX 2023*).
- A. Noroozian, **E. T. Rodríguez**, E. Lastdrager, T. Kasama, M. Van Eeten & C. H. Gañán, “Can ISPs Help Mitigate IoT Malware? A Longitudinal Study of Broadband ISP Security Efforts” *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, 2021, pp. 337-352, doi: 10.1109/EuroSP51992.2021.00031.

DATASETS

Table 8: Datasets available

Publication	Dataset
Rodríguez, E. , Verstegen, S., Noroozian, A., Inoue, D., Kasama, T., van Eeten, M., & Gañán, C. H. (2021). “User compliance and remediation success after IoT malware notifications”. In <i>Journal of Cybersecurity</i> , 7(1), tyab015.	ISP customer data, even anonymized, cannot be shared.
Bouwmeester, B., Rodríguez, E. , Gañán, C., van Eeten, M., & Parkin, S. (2021). “The Thing Doesn’t Have a Name”: Learning from Emergent Real-World Interventions in Smart Home Security. In <i>Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)</i> (pp. 493-512).	ISP customer data, even anonymized, cannot be shared.
E. Rodríguez , M. Fukkink, S. Parkin, M. van Eeten & C. Gañán, “Difficult for Thee, But Not for Me: Measuring the Difficulty and User Experience of Remediating Persistent IoT Malware” <i>2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)</i> , 2022, pp. 392-409, doi: 10.1109/EuroSP53844.2022.00032	ISP customer data, even anonymized, cannot be shared.
Rodríguez, E. , Noroozian, A., van Eeten, M., & Gañán, C. (2021). “Super-Spreaders: Quantifying the Role of IoT Manufacturers in Device Infections”. In <i>20th Annual Workshop on the Economics of Information Security (WEIS 2021)</i> .	Sensitive data, so it cannot be shared beyond the research team.
E. Rodríguez , R. Anghel, S. Parkin, M. van Eeten & C. Gañán, “Two Sides of the Shield: Understanding Protective DNS adoption factors” (<i>USENIX 2023</i>).	Data cannot be shared beyond the research team according to consent forms.

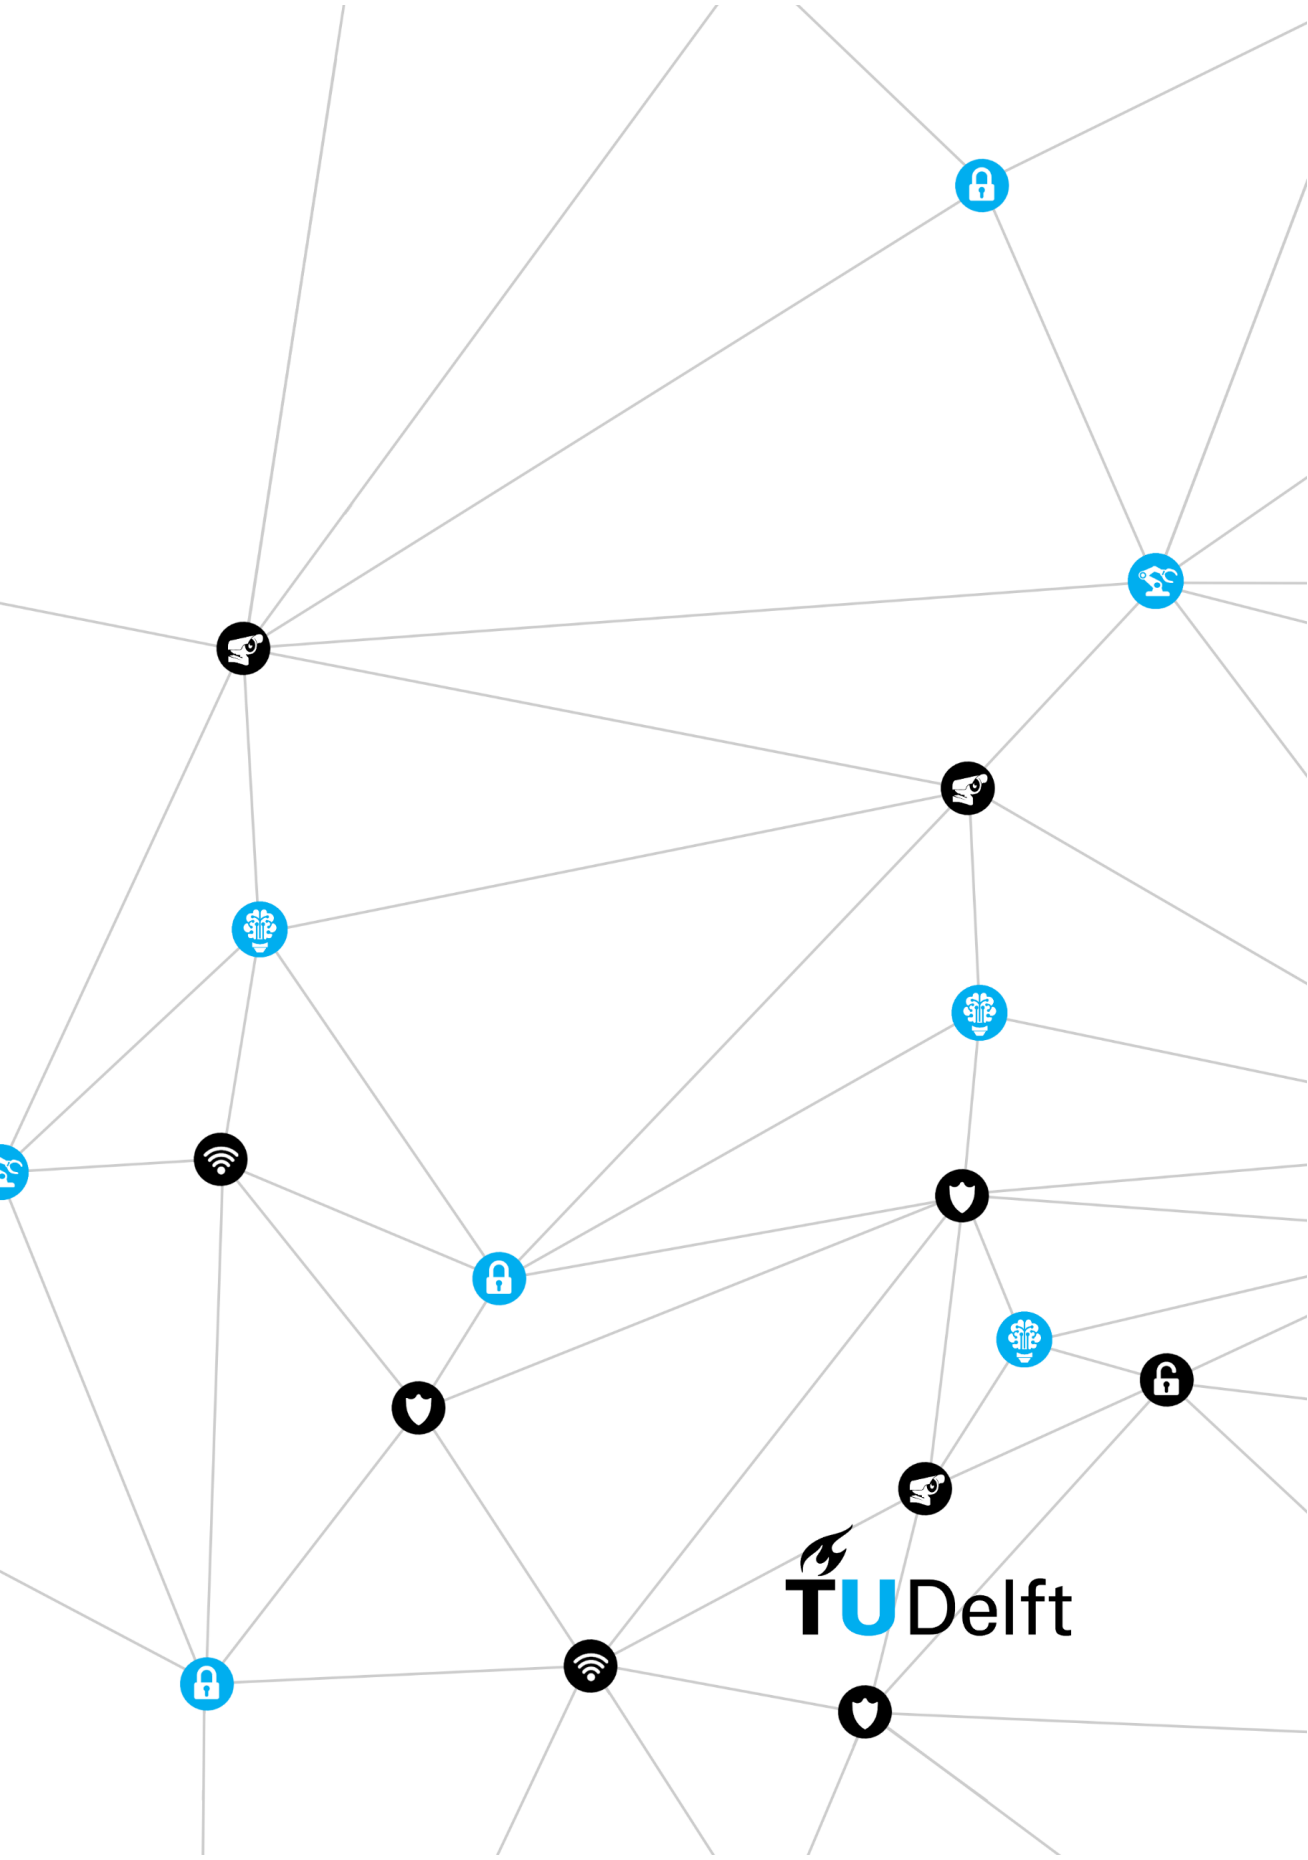
ABOUT THE AUTHOR



Elsa Rebeca Turcios Rodríguez (1986) was born in Masaya, Nicaragua. She received her bachelor's degree (cum laude) from Universidad Iberoamericana de Ciencia y Tecnología in 2008. She worked in the online and content marketing industry and as a Network Engineer back in her home country. In 2016 she moved to The Netherlands, where she pursued her MSc degree in Management of Technology at Delft University. Her master thesis aimed to condense ideas regarding whether E-Privacy Directive approaches were better for privacy protection and present empirical evidence to the privacy discussion to inform policymakers about which legislation features and market factors could help reduce tracking in the E-Privacy Regulation revision. After her graduation, she worked as a consultant for the telecom sector in The Netherlands.

In 2019, she joined the Delft University of Technology as a PhD candidate. Her research was embedded in the Mitigating Internet of Things (IoT)-based distributed denial-of-service (DDoS) via Domain Name System (DNS) project. In this project, she conducted various human-centered studies to understand how users dealt with IoT devices infected with malicious software and the role of manufacturers and intermediaries in supporting users to solve this issue. During this period, she advised five master students until the completion of their graduation projects on related topics.

Elsa's research interests include empirical investigations in security and privacy, economic theories, statistics, and human factors in security.



 **TU**Delft