

## The Signals We Send

### Analysing the Market Signals for IoT Security and Privacy

Vetrivel, Swaathi

**DOI**

[10.4233/uuid:6144bc3f-6101-433c-855a-281fef151862](https://doi.org/10.4233/uuid:6144bc3f-6101-433c-855a-281fef151862)

**Publication date**

2025

**Document Version**

Final published version

**Citation (APA)**

Vetrivel, S. (2025). *The Signals We Send: Analysing the Market Signals for IoT Security and Privacy*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:6144bc3f-6101-433c-855a-281fef151862>

**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.

# The Signals We Send

Analysing the Market Signals for IoT Security and Privacy



Swaathi Vetrivel



# **THE SIGNALS WE SEND:**

**ANALYSING THE MARKET SIGNALS FOR IOT SECURITY AND PRIVACY**



**THE SIGNALS WE SEND:**  
**ANALYSING THE MARKET SIGNALS FOR IOT SECURITY AND PRIVACY**

**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, 28 October, 2025 at 17:30

by

**Swaathi VETRIVEL**

Master of Science in Management of Technology, Delft University of Technology  
born in Chennai, Tamil Nadu, India.

This dissertation has been approved by the promotor:

Prof. dr. M.J.G. van Eeten  
Dr. ir. C. Hernández Gañán

Composition of the doctoral committee:

Rector Magnificus	Chairman
Prof. dr. M.J.G. van Eeten	Delft University of Technology, promotor
Dr. ir. C. Hernández Gañán	Delft University of Technology, promotor

*Independent members:*

Prof. mr. dr. J.A. de Bruijn	Delft University of Technology
Prof. dr. C.P. van Beers	Delft University of Technology
Prof. dr. R. Böhme	University of Innsbruck, Austria
Dr. F.S. Gürses	Delft University of Technology
Dr. A. Abhishta	University of Twente

This research was funded by the RAPID project (Grant No. CS.007), financed by the Dutch Research Council (NWO).



Copyright © 2025 by Swaathi Vetrivel

An electronic version of this dissertation is available at  
<http://repository.tudelft.nl/>

*To Amma and Appa*



# CONTENTS

<b>Summary</b>	<b>xi</b>
<b>Samenvatting</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Security and Privacy Issues with IoT Devices	2
1.2 The Economics of IoT Security and Privacy	3
1.3 Market Signals for Consumers	3
1.4 E-commerce Platforms	4
1.4.1 Actors on E-Commerce Platforms	5
1.5 Research gaps	9
1.6 Research aims and question	10
1.6.1 Study 1 – S&P Signals within Consumer Reviews	12
1.6.2 Study 2 – IoT Device Sales and S&P Signals	12
1.6.3 Study 3 – The Evolution of IoT S&P	12
1.6.4 Study 4 – IoT Update Information on Online Stores	13
1.6.5 Study 5 – Victims of IoT Botnet Attacks	13
<b>2 Consumer Reviews</b>	<b>15</b>
2.1 Introduction	15
2.2 Background and Related Work	17
2.2.1 Analysis of online reviews	17
2.2.2 Consumer perceptions of IoT security and privacy	18
2.2.3 Interventions in consumer IoT purchase decisions	18
2.3 Methodology	19
2.3.1 Selection of IoT devices	19
2.3.2 Product page and review retrieval	20
2.3.3 Review dataset construction	20
2.3.4 Quantifying presence of IoT S&P issues	22
2.3.5 Determining context of IoT S&P themes	23
2.3.6 Research ethics	24
2.4 Results	24
2.4.1 S&P prevalence in reviews	24
2.4.2 Inductive thematic analysis results	26
2.5 Discussion	34
2.5.1 Limitations	35
2.5.2 Recommendations	36
2.6 Conclusions	37

<b>3</b>	<b>IoT Market Dynamics</b>	<b>39</b>
3.1	Introduction	39
3.2	Related Work	42
3.2.1	Consumer purchase decisions and signals for S&P	42
3.2.2	Estimates of IoT device popularity	42
3.2.3	Estimates of security of popular IoT devices	43
3.2.4	Correlation between S&P and sales	43
3.3	Dataset Description	43
3.3.1	Expert IoT device ratings from Consumentenbond	44
3.3.2	Sales data of IoT devices	46
3.3.3	Information from Product Pages	47
3.4	Relationship between S&P and Sales	50
3.4.1	Explanatory Variable Selection	51
3.4.2	Model Evaluation	54
3.4.3	Model Interpretation	55
3.5	Relationship between expert ratings and information from product pages	57
3.5.1	Relationship between expert S&P ratings and average user rating	57
3.5.2	Relationship between update ratings and update support information	58
3.6	Relationship between update information status and sales	59
3.7	Discussion and Recommendations	60
3.7.1	Limitations	61
3.7.2	Research Ethics	62
3.8	Conclusion	62
<b>4</b>	<b>IoT Longitudinal Evolution</b>	<b>63</b>
4.1	Introduction	63
4.2	Related Work	65
4.2.1	IoT Vendor Studies	65
4.2.2	IoT Security Best Practices	66
4.2.3	Security Practices in Organisations	66
4.3	ECA Testing Methodology	66
4.4	Dataset Description	67
4.5	Trends in IoT Security and Privacy Ratings	69
4.5.1	Evolution of S&P Ratings at the Device Type Level	69
4.5.2	Evolution of S&P Ratings at the Manufacturer Level	70
4.6	Trends in IoT Security and Privacy Features	72
4.6.1	Evolution of S&P Features at the Device Type Level	73
4.6.2	Evolution of S&P features at the Manufacturer Level	74
4.7	Factors Influencing IoT Security and Privacy	80
4.8	Discussion	83
4.8.1	Limitations	84
4.9	Conclusion	85

<b>5</b>	<b>IoT Updates</b>	<b>87</b>
5.1	Introduction . . . . .	87
5.2	Related Work . . . . .	89
5.3	Methodology . . . . .	90
5.3.1	Web Store Data Collection . . . . .	91
5.3.2	Evaluating Provisioning: Manual Analysis of Online Stores. . . . .	92
5.3.3	Characterizing Update Durations: Large Scale Analysis . . . . .	92
5.3.4	Comparison Across Sources . . . . .	93
5.4	Provisioning of Update Duration Information on EU Online Stores . . . . .	94
5.4.1	Local Online Stores . . . . .	94
5.4.2	Amazon Stores . . . . .	95
5.4.3	Chinese Online Store - Temu. . . . .	96
5.5	Quantifying Update Support Durations . . . . .	96
5.6	Comparison Across Sources. . . . .	99
5.6.1	Comparison between Dutch Retailers . . . . .	99
5.6.2	Between Retailers and Manufacturer Websites . . . . .	103
5.6.3	Comparison of Retailers, Manufacturers and Database for Smart TVs . . . . .	105
5.7	Discussion . . . . .	107
5.8	Conclusion . . . . .	109
<b>6</b>	<b>Victims of DDoS Attacks</b>	<b>111</b>
6.1	Introduction . . . . .	111
6.2	Background and Related Work . . . . .	114
6.3	Data Sources and Methodology . . . . .	116
6.3.1	NetLab Data . . . . .	116
6.3.2	AmpPot Data. . . . .	117
6.4	Methodology . . . . .	118
6.5	Results . . . . .	121
6.5.1	Distribution of targeted AS types. . . . .	121
6.5.2	Rankings of the targeted ASes . . . . .	122
6.5.3	Geographical distribution of the targeted IP addresses. . . . .	123
6.5.4	Comparison of domains hosted in the targeted IPs . . . . .	124
6.5.5	Analysis of domains resolving to top 100 most common IPs. . . . .	126
6.5.6	Duration of the DDoS attacks . . . . .	128
6.5.7	Modelling . . . . .	129
6.6	Discussion . . . . .	130
6.7	Conclusion . . . . .	134
<b>7</b>	<b>Conclusion</b>	<b>135</b>
7.1	Discussion . . . . .	137
7.1.1	Consumers. . . . .	137
7.1.2	Manufacturers . . . . .	139
7.1.3	E-commerce platforms . . . . .	139
7.1.4	Societal Implications. . . . .	140

7.2	Implications for Governance and Policy Making . . . . .	141
7.2.1	Hierarchical Governance. . . . .	141
7.2.2	Network Governance. . . . .	142
7.2.3	Market Governance . . . . .	143
7.3	Future Work. . . . .	143
	<b>Bibliography</b>	<b>145</b>
<b>A</b>	<b>Appendix</b>	<b>167</b>
A.1	Search terms used for each of the four device types . . . . .	167
A.2	Search terms used for identifying S&P related customer reviews by category	167
A.3	Results of LDA . . . . .	167
A.4	Results of Anchored CorEx . . . . .	168
A.5	Codebook from Inductive Coding. . . . .	168
<b>B</b>	<b>Appendix</b>	<b>171</b>
B.1	List of sub-tests for each device type . . . . .	171
B.2	Price comparison between CB tested and non-tested IoT devices . . . . .	172
B.3	GLM Variable Selection . . . . .	172
B.4	Correlation between CB device ratings and average ratings on Amazon and Winkel . . . . .	173
B.5	GLM Results with Update Statuses . . . . .	173
<b>C</b>	<b>Appendix</b>	<b>175</b>
C.1	Beta Regression Estimates . . . . .	176
	<b>Acknowledgements</b>	<b>177</b>
	<b>Authorship Contributions</b>	<b>181</b>
	<b>List of Publications</b>	<b>183</b>
	<b>About the Author</b>	<b>185</b>

# SUMMARY

The rapid rise in Internet-of-Things (IoT) devices, from smart thermostats and fitness trackers to connected cameras, while providing unprecedented convenience to consumers and profitable subscription based business models to manufacturers, has also raised critical security and privacy (S&P) concerns. From hacked video feeds and exploitation of sensitive data to revenue loss from service outages due to Distributed Denial of Service (DDoS) attacks, the consequences of poor S&P of IoT devices are experienced both at the individual level and at the collective societal level.

The underlying reasons for the S&P issues in IoT devices are not merely technical, there are socio-technical and economic dimensions associated with them. For instance, large scale DDoS attacks from insecure IoT devices are a classic example of negative externalities where the consequences of the attack are experienced by a party that is neither the manufacturer nor the consumer. In such a context, manufacturers often face a lack of incentives to improve on the underlying S&P issues since doing so would increase their development costs and delay their time to market. Although consumers as device owners may not be directly targeted by DDoS attacks, they do face indirect consequences from DDoS attacks on governments, banks and other websites. Moreover, they bear the brunt of individual losses to S&P, for example, when their IoT devices are hacked or their personal video feeds are exposed. Therefore, consumers have incentives to buy IoT devices with strong S&P features. Recent studies affirm this, and show that consumers not only care about IoT S&P, they are also willing to pay a premium for it – if they are informed about the S&P at the time of purchase.

However, the problem still remains that consumers do not have sufficient information – at the time of purchase – to discern IoT devices that have good S&P features from those that do not. While regulations like the Cyber Resilience Act (CRA) in the EU, and the US Cyber Trust Mark aim to decrease this information asymmetry, they are not yet in effect. In the absence of official information about an IoT device's S&P at the time of purchase, consumers might use other signals that directly or indirectly indicate the S&P posture of IoT devices like mention of security concerns in consumer reviews on e-commerce platforms. Since consumers currently depend on such indirect sources to assess S&P, insights into these signals can help design more effective interventions that fit into their current decision making flow. However, there is currently no empirical analysis on these market signals which limits our understanding of how much the consumer base already recognises and signals a need for S&P.

This dissertation addresses this gap by analyzing S&P of consumer IoT devices through a market-based empirical lens that examines how economic incentives, S&P signals, and purchase decisions interact across different stakeholders in real-world e-commerce settings. Specifically, five mature and popular IoT device types are considered: IP cameras, smart printers, smart speakers, smart TVs and smart watches. By examining the interactions between manufacturers, consumers, sellers, and the e-commerce platforms that

sell these devices, using actual market data (sales figures, prices, reviews, and product listings), this dissertation provides a unique vantage point on the market signals for IoT S&P and information asymmetry experienced by consumers. Overall, this dissertation aims to answer the following overarching research question through five research studies.

*What signals for security and privacy are present in the e-commerce platforms that sell IoT devices?*

The five studies in this dissertation together create a comprehensive picture of the IoT S&P market ecosystem that no single study could provide. In chapter 2, we investigate to what extent customer reviews of IoT products provide S&P information to consumers at the point of purchase. We found that around 10% of consumer reviews contains S&P signals which included technical statements about features, frustrations with specific device use activities, as well as vignettes about trying to use a device in a particular context. Negative views on IoT devices were reflected in generally lower overall ratings for devices. All in all, we found that customer reviews provide a practical and widely used mechanism for conveying S&P information to consumers — a valuable complement to potential future labelling schemes for IoT. We recommended mechanisms to amplify these signals within the e-commerce platforms to encourage consumers to make more S&P conscious purchase decisions.

Our analysis on whether consumer preferences for S&P are reflected in their real life purchase decisions is presented in Chapter 3. We studied the correspondence between the sales of IoT devices and the signals for S&P in the e-commerce platforms. We constructed a model with the device sales as the dependent variable and price, expert S&P rating and overall rating, user ratings and review count from two e-commerce platforms as the predictor variables. Our results showed that despite lack of information about S&P for a majority of the devices, a one standard deviation increase in the S&P rating of a device was associated with a 56% increase in sales. However, we also observed that the relationship was moderated by the price of the device. The effect was stronger at lower prices and decreased as price increased. Further, we also find that on one of the online stores, devices with complete update support information correspond to a 69% increase in sales. This suggests that there are positive incentives at play that will reward manufacturers and sellers who adopt S&P transparency initiatives like security labels with higher sales. Our results also highlighted the crucial role of sellers in ensuring the success of such initiatives since they are responsible for updating the relevant information on online stores.

We present our study on the evolution of S&P features in IoT devices over the past decade in Chapter 4. We studied 428 IoT devices from 23 manufacturers, focusing on three widely used and mature IoT device types: IP cameras, smart printers, and smart speakers. Our findings showed that while IP cameras maintained consistently high S&P ratings, smart printers exhibited lower ratings with a slight declining trend, and smart speakers had the lowest ratings with no clear temporal pattern. At the manufacturer level, only a minority (3 out of 23) demonstrated improvement, while the majority (12 out of 23) maintained stable S&P ratings, and the remaining eight showed no clear trend in the ratings. Our analysis also uncovered a surprising trend of inconsistent deployment of S&P features in subsequent device models of the same manufacturer. We find that stability in ratings obscured such underlying inconsistencies. This highlights a need to

help manufacturers operationalise S&P best practices within their development processes to ensure a more systematic and coherent focus on S&P.

In Chapter 5, we examine the availability of update support information for IoT devices on e-commerce platforms in the EEA countries. We analysed over 26,000 product listings across major online retailers, including regional EU stores, Amazon, and Temu, as well as manufacturer websites and – for smart TVs – the centralised EU EPREL database. We identified significant gaps and inconsistencies in the information available to consumers. However, our findings also show that regulatory measures and targeted transparency interventions are associated with better information disclosure. Dutch e-commerce platforms, which were subject to a policy intervention for update support disclosure, had high rates of availability, while platforms like Temu and Amazon disclosed little to none. The centralised EU energy labelling database, which contains update support information alongside energy ratings for electronic displays such as smart TVs, had the highest percentage of availability, likely due to stronger enforcement mechanisms. We also observed discrepancies in the update duration across different sources raising concerns about consistency and accuracy. For the same device, the update duration on the e-commerce platform differed from that on the manufacturer websites and, for some smart TVs, also from the duration stated in the centralised energy database. Our findings highlight both the potential and the limitations of current regulatory efforts, and underscore the need for mechanisms to ensure not just disclosure, but also accuracy.

The negative externalities that arise from a lack of strong S&P measures in IoT devices, specifically in the context of IoT botnets, are the focus of Chapter 6. Our research showed that the commoditization of IoT botnets may alter the technological supply in the DDoS booter services. We did not observe repeated victimisation to increase over the years, meaning that the increase in frequency of DDoS attacks is tied to a growth in the number of victims. However, we see that the newer targets operate in large variety of sectors, with small medium enterprises getting more attacks and gaming services still being a common target. This illustrates the broader negative externalities of DDoS attacks and stresses the urgency to adopt stronger legal actions against miscreants launching DDoS attacks.

Overall, this dissertation analyses the market signals sent and used by different actors in the e-commerce ecosystem — consumers, manufacturers, sellers, and the platforms themselves — and examines the broader consequences of market failures, from information asymmetry at the point of purchase to large-scale DDoS attacks exploiting insecure devices. We show that S&P signals and information, even when available, are fragmented, inconsistent, and often invisible at the critical moment of purchase. On their part, consumers express their experiences and concerns about S&P of IoT devices organically in e-commerce platforms both directly, through consumer reviews, and indirectly through their purchase choice. What is still missing is a mechanism to highlight these concerns to manufacturers so they can act on them. While manufacturers invest effort in S&P, the inconsistent deployment of S&P features across subsequent device models from the same manufacturer underscores a need to streamline their development processes. The e-commerce platform design also plays a crucial role in encouraging more S&P conscious purchase decisions. Ultimately, all of these actors must work together to ensure that IoT devices with strong S&P features enter the market, that S&P information is available at the point of purchase, and that consumers can make informed decisions.



# SAMENVATTING

De snelle toename van Internet-of-Things (IoT)-apparaten – van slimme thermostaten en fitnesstrackers tot verbonden camera's – biedt consumenten ongekende gemakken en fabrikanten winstgevende abonnementsmodellen, maar brengt ook kritieke zorgen met zich mee over cybersecurity en privacy. Van gehackte videobeelden en misbruik van gevoelige gegevens tot omzetverlies door serviceonderbrekingen als gevolg van Distributed Denial of Service (DDoS)-aanvallen: de gevolgen van gebrekkige cybersecurity en privacy van IoT-apparaten worden zowel op individueel niveau als op collectief maatschappelijk niveau ervaren.

De onderliggende oorzaken van de cybersecurity- en privacyproblemen bij IoT-apparaten zijn niet louter technisch, maar hebben ook sociaal-technische en economische dimensies. Grootschalige DDoS-aanvallen vanuit onveilige IoT-apparaten zijn bijvoorbeeld een klassiek voorbeeld van negatieve externaliteiten, waarbij de gevolgen worden ervaren door partijen die noch de fabrikant, noch de consument zijn. In zo'n context hebben fabrikanten vaak weinig prikkels om de onderliggende cybersecurity- en privacyproblemen aan te pakken, omdat dit hun productiekosten zou verhogen en hun time-to-market zou vertragen. Hoewel consumenten die in het bezit zijn van slimme apparaten niet rechtstreeks worden getroffen door DDoS-aanvallen, ondervinden ze wel indirecte gevolgen van aanvallen op overheden, banken en andere websites. Bovendien ondervinden zij de gevolgen van individuele incidenten op het gebied van cybersecurity en privacy, bijvoorbeeld wanneer hun IoT-apparaten worden gehackt of hun persoonlijke videobeelden openbaar worden. Consumenten kunnen daardoor eerder gemotiveerd zijn om IoT-apparaten te kopen met sterke cybersecurity- en privacy-eigenschappen. Recente studies bevestigen deze aanname en tonen aan dat consumenten niet alleen geven om IoT-cybersecurity en -privacy, maar ook bereid zijn er extra voor te betalen – mits ze op het moment van aankoop geïnformeerd worden over de cybersecurity en privacy.

Het probleem blijft echter dat consumenten op het moment van aankoop niet over voldoende informatie beschikken om te onderscheiden welke IoT-apparaten goede cybersecurity- en privacy-eigenschappen hebben en welke niet. Hoewel regelgeving zoals de Cyber Resilience Act (CRA) in de EU en het Amerikaanse Cyber Trust Mark bedoeld zijn om deze informatie-asymmetrie te verkleinen, zijn deze nog niet van kracht. Bij gebrek aan officiële informatie over de cybersecurity en privacy van een IoT-apparaat op het moment van aankoop kunnen consumenten andere signalen gebruiken die direct of indirect de cybersecurity- en privacy-positie van IoT-apparaten aangeven, zoals meldingen van beveiligingsproblemen in consumentenreviews op e-commerceplatforms. Aangezien consumenten momenteel afhankelijk zijn van zulke indirecte bronnen om cybersecurity en privacy te beoordelen, kunnen inzichten in deze signalen helpen om effectievere interventies te ontwerpen die passen bij hun huidige besluitvormingsproces. Er is echter nog geen empirische analyse van deze marktsignalen, waardoor we nog onvoldoende begrijpen in hoeverre consumenten al een behoefte aan cybersecurity en

privacy herkennen en uiten.

Dit proefschrift draagt bij aan het invullen van deze leemte door de cybersecurity en privacy van consumentgerichte slimme apparaten te analyseren vanuit een empirisch marktgericht perspectief, met bijzondere aandacht voor de manier waarop economische prikkels, signalen en aankoopbeslissingen elkaar beïnvloeden tussen verschillende belanghebbenden in reële e-commerceomgevingen. Specifiek worden vijf gangbare en populaire typen IoT-apparaten beschouwd: IP-camera's, slimme printers, slimme luidsprekers, smart-tv's en smartwatches. Dit proefschrift onderzoekt de interacties tussen fabrikanten, consumenten, verkopers en de e-commerceplatforms waarop deze apparaten worden verkocht – aan de hand van reële marktgegevens (verkoopcijfers, prijzen, reviews en productvermeldingen) – en biedt een uniek perspectief op de marktsignalen rond beveiliging en privacy van IoT-apparaten en de informatie-asymmetrie waarmee consumenten geconfronteerd worden. In het algemeen beoogt dit proefschrift de volgende overkoepelende onderzoeksvraag te beantwoorden via vijf deelstudies:

*Welke signalen voor veiligheid en privacy zijn aanwezig op de e-commerceplatforms die IoT-apparaten verkopen?*

De vijf studies in dit proefschrift vullen elkaar aan en bieden samen een rijker inzicht in het IoT-cybersecurity- en privacy-marktecosysteem dat geen enkele afzonderlijke studie alleen zou kunnen bieden. In hoofdstuk 2 onderzoeken we in welke mate klantreviews van IoT-producten cybersecurity- en privacy-informatie aan consumenten verschaffen op het moment van aankoop. We ontdekten dat ongeveer 10% van de consumentenreviews cybersecurity- en privacy-signalen bevatte, waaronder technische opmerkingen over functies, frustratie bij specifieke gebruikssituaties en korte beschrijvingen van pogingen om een apparaat in een bepaalde context te gebruiken. Negatieve opvattingen over IoT-apparaten kwamen tot uiting in lagere algemene waarderingen. Al met al concluderen we dat klantreviews een waardevol en veelgebruikt mechanisme vormen om cybersecurity- en privacy-informatie aan consumenten over te brengen – voorafgaand aan en aanvullend op toekomstige etiketteringsregelingen voor IoT. We bevelen aan om mechanismen te ontwikkelen die deze signalen versterken binnen e-commerceplatforms om consumenten te stimuleren meer cybersecurity- en privacy-bewuste aankoopbeslissingen te nemen.

Onze analyse of consumentenvoorkeuren voor cybersecurity en privacy zich ook daadwerkelijk vertalen in aankoopgedrag wordt gepresenteerd in hoofdstuk 3. We onderzochten het verband tussen de verkoop van IoT-apparaten en de cybersecurity- en privacy-signalen op e-commerceplatforms. We stelden een model op met de verkoopcijfers als afhankelijke variabele en prijs, deskundige cybersecurity- en privacy-beoordeling, algemene beoordeling, gebruikersbeoordelingen en aantal reviews van twee e-commerceplatforms als verklarende variabelen. Onze resultaten toonden aan dat, ondanks het gebrek aan informatie over cybersecurity en privacy voor de meeste apparaten, een stijging van één standaarddeviatie in de cybersecurity- en privacy-beoordeling gepaard ging met een stijging van 56% in de verkoop. We zagen echter ook dat deze relatie werd beïnvloed door de prijs van het apparaat: het effect was sterker bij lagere prijzen en nam af bij hogere prijzen. Verder bleek dat op één van de webwinkels apparaten waarvoor volledige informatie over updatesupport beschikbaar was, 69% hogere verkoopcijfers hadden. Dit suggereert dat er positieve prikkels bestaan die fabrikanten en verkopers be-

lonen voor het aannemen van transparantie-initiatieven met betrekking tot cybersecurity en privacy, zoals beveiligingslabels. Onze resultaten benadrukken ook de cruciale rol van verkopers bij het succes van dergelijke initiatieven, aangezien zij verantwoordelijk zijn voor het bijwerken van de relevante informatie in online winkels.

In hoofdstuk 4 presenteren we ons onderzoek naar de evolutie van cybersecurity- en privacy-functies in IoT-apparaten gedurende het afgelopen decennium. We onderzochten 428 IoT-apparaten van 23 fabrikanten, met de focus op drie veelgebruikte en gangbare typen: IP-camera's, slimme printers en slimme luidsprekers. Onze bevindingen tonen aan dat IP-camera's consequent hoge cybersecurity- en privacy-beoordelingen behielden, slimme printers lagere beoordelingen kregen met een licht dalende trend, en slimme luidsprekers de laagste scores hadden zonder duidelijk patroon. Op fabrikantsniveau lieten slechts drie van de 23 fabrikanten verbeteringen zien; twaalf bleven stabiel en acht vertoonden geen duidelijk patroon. Onze analyse bracht ook een verrassende trend aan het licht: de inconsistente implementatie van cybersecurity- en privacy-functies in opeenvolgende modellen van dezelfde fabrikant. Deze schijnbare stabiliteit in beoordelingen maskeerde dergelijke onderliggende inconsistenties. Dit benadrukt de noodzaak om fabrikanten te helpen cybersecurity- en privacy-best practices te operationaliseren binnen hun ontwikkelprocessen om een systematischere en coherente focus op cybersecurity en privacy te waarborgen.

In hoofdstuk 5 onderzoeken we de beschikbaarheid van informatie over updatesupport voor IoT-apparaten op e-commerceplatforms in de EER-landen. We analyseerden meer dan 26.000 productvermeldingen bij grote online retailers, waaronder regionale EU-winkels, Amazon en Temu, evenals websites van fabrikanten en – voor smart-tv's – de gecentraliseerde EU EPREL-database. We identificeerden aanzienlijke hiaten en inconsistenties in de informatie die beschikbaar is voor consumenten. Onze bevindingen tonen aan dat regelgeving en gerichte transparantie-interventies samenhangen met betere informatieverstrekking. Nederlandse e-commerceplatforms, die onderhevig waren aan een beleidsinterventie omtrent updatesupport, hadden hoge beschikbaarheidspercentages, terwijl platforms zoals Temu en Amazon weinig tot geen informatie hierover vermeldden. De gecentraliseerde EU-database voor energie-etikettering, die naast energiebeoordelingen ook informatie over updatesupport bevat voor elektronische displays zoals smart-tv's, had het hoogste beschikbaarheidspercentage, waarschijnlijk door een strengere handhaving. We observeerden ook discrepanties in de duur van updates tussen verschillende bronnen, wat vragen oproept over consistentie en nauwkeurigheid. Voor hetzelfde apparaat verschilde de updatesupportduur op het e-commerceplatform van die op de website van de fabrikant en, bij sommige smart-tv's, ook van die in de centrale energiedatabase. Onze bevindingen benadrukken zowel de mogelijkheden als de beperkingen van huidige regelgeving en onderstrepen de noodzaak van mechanismen die niet alleen openbaarmaking maar ook juistheid waarborgen.

De negatieve externaliteiten die voortvloeien uit een gebrek aan sterke cybersecurity- en privacy-maatregelen in IoT-apparaten, met name in de context van IoT-botnets (waarbij een botnet een netwerk is van geïnfecteerde apparaten die op afstand worden bestuurd), vormen de kern van hoofdstuk 6. Ons onderzoek toont aan dat de commercialisering van IoT-botnets het technologische aanbod binnen DDoS-booterdiensten kan beïnvloeden. We zagen geen toename in herhaalde slachtofferschap door de jaren heen,

wat suggereert dat de toename in DDoS-aanvallen eerder samenhangt met een groeiend aantal slachtoffers. We constateren dat de nieuwe doelwitten actief zijn in een breed scala aan sectoren, waarbij kleine en middelgrote ondernemingen vaker worden getroffen en gamingdiensten nog steeds een veelvoorkomend doelwit vormen. Dit illustreert de bredere negatieve externaliteiten van DDoS-aanvallen en benadrukt de urgentie van strengere juridische maatregelen tegen daders van dergelijke aanvallen.

# 1

## INTRODUCTION

In recent years, there has been a rapid proliferation in the number of internet connected 'things' – yesteryear things have now become 'smart devices' through added functionality enabled by internet connectivity. There are endless examples of these internet connected devices, from TVs and washing machines to coffee machines and thermostats. These devices are referred to as Internet-of-Things (IoT), and defined by the ITU as 'a global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies' [2]. At the end of 2023, there were 16.6 billion IoT devices worldwide, more than double the human population [161, 230].

The adoption rate of IoT devices will undoubtedly continue to increase in the coming years – the number of IoT devices is projected to reach 40 billion by 2030 [107]. The explosive increase in the adoption of IoT devices is an artefact of both a demand pull and supply push. One of the main reasons for the increase in consumer demand are the conveniences and functionalities enabled by them. For example, internet connected cameras allow users to monitor their homes remotely providing real time alerts on intrusion detection. Smartwatches and fitness trackers provide continuous analytics based on vitals increasing user awareness of health and enabling early detection of anomalies. Smart thermostats can automatically regulate the indoor temperature based on user preferences and weather patterns and also allow for remote operation – you can now turn on your heater from your phone on the flight back from your tropical vacation and be welcomed into a warm and cosy home. The supply push is also evident since these services turn one off purchases into continuous subscriptions creating new business models for revenue growth.

The benefits of IoT devices, however, come with significant Security and Privacy (S&P) risks. Internet connected cameras can expose personal video feeds on public internet, sensitive health data from smart watches can be exploited for manipulative marketing and attackers can hijack smart thermostats and control home heating systems. As IoT devices become more pervasive and deeply embedded into our daily life and critical functionalities like heating, ensuring their security and the protection of user privacy

becomes crucial.

### 1.1. SECURITY AND PRIVACY ISSUES WITH IoT DEVICES

There have been numerous instances that have raised concerns about the lack of adequate S&P controls in IoT devices. For instance, hackers took control of Ring cameras – commonly used for home monitoring – and exploited its two way communication feature to interact with children and families in disturbing ways [215]. Due to a technical issue, users of Xiaomi cameras were able to view random video feeds of other users [212]. Studies have shown that information collected via smart meters can be used to deduce details about the occupants of the house, including the number, gender and religion [140]. Other risks include but are not limited to unauthorised access and misuse of sensitive data, substantial risks to personal safety due to disruption of critical services and lack of update provisioning leaving the devices vulnerable to even known security flaws [83]. These individual losses represent direct harm to consumers who suffer privacy violations, financial losses, or safety risks from compromised devices. However, the consequences of inadequate IoT S&P extend far beyond individual users to create externalities at the societal scale.

For instance, the large number of IoT devices has made it an attractive target for attackers who compromise these devices at scale and recruit them into IoT botnets. A botnet is a collection of internet-connected compromised devices known as 'bots' that can be centrally controlled by an attacker. These botnets are typically used to launch Distributed Denial of Service (DDoS) attacks that overwhelm a target web service or Internet infrastructure with malicious traffic rendering it incapable of processing legitimate requests [31].

The Mirai malware [124] is a notorious example of the large-scale havoc that can be caused by IoT botnets. The attack vector of Mirai was relatively simple. It scanned large portions of the IPv4 address space for devices with open Telnet ports and attempted to log in using common default username-password combinations (e.g., admin:admin). Upon successful authentication, the malware installed and executed the Mirai binary on the device, thereby infecting it. The infected device would then connect to a central server to receive instructions and begin scanning the internet for other vulnerable hosts, continuing the propagation cycle.

In October 2016, the Mirai botnet was responsible for a massive DDoS attack against the DNS provider Dyn. This attack, which peaked at a reported 1.2 Tbps, temporarily disrupted access to major websites including GitHub, Twitter, Reddit, Netflix, and Airbnb. Mirai marked a turning point in the evolution of botnets. The simplicity of its attack vector and its rapid spread demonstrated how a large number of poorly secured IoT devices could be weaponized to compromise even well-defended targets [23]. More recently, in 2025, the popular cybersecurity blog KrebsOnSecurity was hit with a DDoS attack of 6.3 Tbps by an IoT botnet named Aisuru/Airashi, from the same Mirai family [128]. Although the attack lasted only 45 seconds and was possibly a test run to demonstrate the fire power of the new botnet, very few websites are currently equipped to handle DDoS attacks at this scale. While DDoS attacks predate the proliferation of IoT, the sheer volume and diversity of IoT devices have made them especially attractive to attackers, enabling attacks of unprecedented scale and magnitude. In addition to DDoS attacks, IoT botnets

have also been used to mine cryptocurrencies for industrial espionage and to steal online banking credentials [123].

## 1.2. THE ECONOMICS OF IOT SECURITY AND PRIVACY

DDoS attacks and other problems caused by IoT botnets are a special case of security failures due to IoT devices since the adverse consequences are experienced by third parties – such as DNS providers – that are neither the user, the seller nor the manufacturer of these devices. This represents a classic negative externality: the costs and consequences of security failures are borne by parties who had no role in the original transaction or device design decisions. When those responsible for protecting systems do not bear the full costs of security failures, market failures emerge, leading to systematic underinvestment in S&P. In economic terms, this misalignment of costs and consequences leads to inefficient allocation of resources and a net loss of economic value, such as the business disruptions and financial losses associated with large-scale DDoS attacks [21].

From a traditional economic perspective, negative externalities such as DDoS attacks are seen as a consequence of misaligned incentives. The parties that suffer the consequences are the victims of the attacks, the network operators and the service providers, not the consumers or the manufacturers. This limits the incentives for manufacturers to improve the S&P of IoT devices. Furthermore, in the context of DDoS attacks from IoT botnets, it is not yet clear where the negative externalities end up or who the victims of IoT botnet attacks are. Prior to the advent of IoT botnets, victims of DDoS attacks were mostly in broadband access networks, with relatively fewer targets in hosting and enterprise networks [170]. However, it is still unclear whether this victimisation pattern holds for IoT botnet-based DDoS attacks. A better understanding of the victimisation pattern would help in designing targeted interventions and defence measures.

## 1.3. MARKET SIGNALS FOR CONSUMERS

Although consumers as device owners are not directly targeted by DDoS attacks, they do face indirect consequences from DDoS attacks on governments, banks and other websites. Moreover, they bear the brunt of individual losses to S&P, for example, when their IoT devices are hacked or their personal video feeds are exposed. Therefore, consumers might have incentives to buy IoT devices with strong S&P features. Recent studies affirm this, and show that consumers not only care about IoT S&P, they are also willing to pay a premium for it if they are informed about the S&P at the time of purchase [34, 70, 89].

Yet a critical question remains: *can consumers actually act on these preferences?* Willingness to pay is only meaningful if consumers have sufficient information to distinguish secure devices from insecure ones at the moment of purchase. In the current market, however, there is scarce information available to consumers about the S&P posture of IoT devices [156]. This creates what economists call information asymmetry or a ‘market for lemons’ [14], where sellers have more information about product quality than buyers, potentially leading to a race to the bottom where poor-quality products drive out good ones. Although there are initiatives in different countries that aim to improve transparency – from the Cyber Trust Mark in the US [79] to the Product Security and Telecommunications Infrastructure (PSTI) in the UK [66] and the Cyber Resilience Act (CRA) in the European

Union [74] – they are still in infancy and not yet implemented fully.

In the absence of official information about an IoT device's S&P at the time of purchase, consumers might use other signals that directly or indirectly indicate the S&P posture of IoT devices. For instance, references to devices being hacked in customer reviews on e-commerce platforms like Amazon could be taken as a signal of poor S&P. However, there is currently no research on market signals in the context of IoT S&P. Understanding these organic market signals is crucial for several reasons. First, consumers are already making purchase decisions today based on whatever information is available – understanding what signals they use and how reliable these signals are can help design more effective interventions. Second, even when official labelling schemes are implemented, they will exist alongside these organic signals; understanding their interplay is essential for policy design. Third, if certain types of information – such as mention of security incidents in reviews – already influence purchase decisions, this creates immediate incentives for manufacturers to improve S&P even before formal regulations take effect.

Prior research has demonstrated that consumers do respond to various market signals in their purchase decisions. Martín and Camarero [147] show that consumers develop trust in firms through consistent signals, such as warranty, advertising or customer service, sent by them. Similarly, when a brand sends transparency signals through public disclosures or statements, there is an increase in perceived brand integrity which influences downstream behavioural intentions [44]. An examination of the effects of marketing and non-marketing controlled signals showed that non-marketing controlled, third party signals such as quality ratings reduce the effectiveness of pricing and advertising while enhancing the credibility of warranties [13]. In online markets, Kristin et al. [129] find that consumer reviews, price, and consumer trust are crucial factors that influence purchase decisions. Moreover, it has been shown that the emotional content of consumer reviews has a positive influence on product quality and increases purchase likelihood [227]. At the same time, consumers react to price signals differently based on context and involvement. Dutta and Bhowmick [68] show that consumers challenge claims on lowest price guarantees especially in online platforms, when they spend some cognitive elaboration cycles on it.

However, the lack of empirical analysis on S&P-specific market signals in the IoT context limits our understanding of how much the consumer base already recognises and expresses a need for S&P, and whether this recognition translates into purchase behaviour that rewards manufacturers of more secure devices. This dissertation addresses this gap by analysing the S&P related market signals currently present on e-commerce platforms and examining whether they influence actual purchase outcomes.

## 1.4. E-COMMERCE PLATFORMS

We focus on e-commerce platforms or online retail stores because they are rapidly becoming the dominant channels for consumer purchases, including for the purchase of IoT devices. In 2024, global sales through e-commerce platforms reached 6.09 trillion USD, marking an 8.4% increase from 2023 [202]. As the dominance of e-commerce platforms increases, they also serve as the primary interface through which information about device S&P is presented and the context within which consumers make purchase decisions. Moreover, omni-channel purchases are also gaining popularity, consumers

explore e-commerce stores and purchase in physical stores or vice versa.

The integration of online and offline retail channels highlights how their distinct features serve complementary roles in consumer purchases. While both facilitate consumer purchases, online stores have a few distinct features that differ from the traditional retail stores. These features are often designed to mimic or substitute for elements that are inherently present in physical stores but absent online, such as in-person assistance, the ability to physically inspect products or other subtle reputation or quality indicators like store location or presentation. For instance, e-commerce platforms allow consumers to rate and review products, which serves as a digital proxy for quality assessment. Just as a centrally located physical store may signal higher quality or reliability, products with high ratings in e-commerce platforms may be perceived more favourably [191]. Similarly, customer reviews and detailed product descriptions serve as a substitute for in-store service, enabling prospective buyers to learn more about the product from the seller provided descriptions and from the experiences of others in the absence of direct interaction with sales staff.

Reviews and other elements of an online store may also serve as indirect indicators of IoT device quality including its S&P posture. Online product reviews have been recognised as containing critical information regarding consumers' concerns [134] and are considered essential in building a firm's business intelligence [236]. Reviews have been noted to function both as informants and as recommenders influencing both product sales and purchase decisions [176]. Studies on the impact of online reviews show that the quality and quantity of product reviews positively influence the sales of a product [139]. Prior studies have also analysed online reviews of certain types of IoT devices like smart home assistants and wearables to understand if there are any privacy concerns. The results vary between the studies from a significant percentage of users concerned with privacy [146] to only 2% [82].

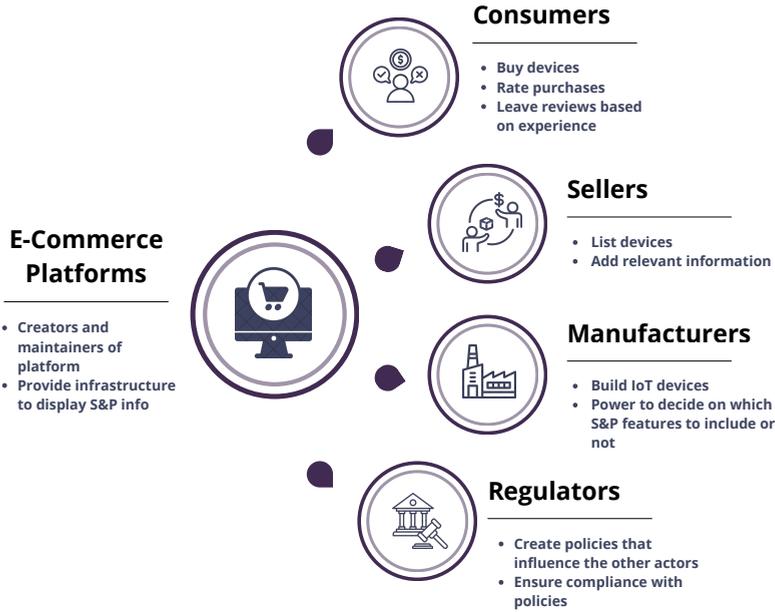
#### 1.4.1. ACTORS ON E-COMMERCE PLATFORMS

To better understand how various elements within e-commerce platforms might signal (or fail to signal) the S&P posture of IoT devices at the time of purchase, this section examines the key actors involved in these platforms and the actions enabled by the platform design. Analysing these actors – consumers, sellers, manufacturers and regulators – alongside the role played by the e-commerce platforms themselves, can help understand the signals they send and how their interactions shape consumer perceptions of S&P at the point of purchase.

##### CONSUMERS

Consumers play a central role in e-commerce platforms, and have the power to create demand for IoT devices with better S&P through their purchase decisions. Prior interviews with retail customers have highlighted the point of purchase as a critical point for informing decisions about the security of new computing devices [177]. However, as discussed earlier, there is often insufficient information about S&P of IoT devices on online stores. To reduce this information asymmetry, research has advocated for IoT security standards and labels [20, 39, 118].

Various forms of S&P information labels have been proposed, including graded label,



**Figure 1.1:** Actors and their actions on E-Commerce Platforms

labels outlining S&P features, and labels indicating ‘approval’ by independent assessment. Emami-Naeini et al. [72] presented participants with labels which were a mix of these label types and report positive feedback from participants who also indicated that they struggled to find this kind of information at the time of device purchase. Johnson et al. [111] found that except for cases where an information label indicated that a device has poor security, consumers were significantly more likely to buy a device with a label. They also note that although functionality was generally more valued, people were willing to pay the same premium for both improved functionality and security.

Apart from labels, prior research has explored other methods and mediums to inform consumers of the privacy posture of devices and their impact on purchase decisions. For instance, Gopavaram et al. [90] investigated customers’ Willingness-To-Pay (WTP) for privacy vs their Willingness-To-Accept (WTA) a lack thereof through an emulated marketplace study. WTA participants, presented with the highest privacy settings by default, were more likely to pay a premium and purchase devices with a higher privacy rating, thereby indicating that interface design also influences purchase decisions.

With a focus on the availability and regularity of security updates for IoT devices, Morgner et al. [157] found support among their survey participants for update related information, more so for those who perceived higher risks in using such devices. Along similar lines, Blythe et al. [34] note that providing people with simple security-related information prior to making their purchase decision has the potential to encourage the purchase of devices which are more secure. They also found that consumers are willing to pay more for increased security and that the relative amount of risk reduction has no significant impact on that willingness.

In contrast, a study by Williams et al. [229] observes that the price of consumer IoT devices deterred more users from buying than privacy concerns and note that since the purchase of IoT devices is voluntary, privacy was more likely to be sacrificed for functionality rather than necessity. This trade-off is echoed in other studies [96, 239], which observe that users balance the risks of using IoT devices against the convenience and benefits offered.

Customer reviews of other devices, including certain types of IoT devices like smart home assistants and wearables, have been studied previously. Two such studies [82, 146] use unsupervised machine learning techniques to understand if users express any privacy concerns in reviews of popular e-commerce sites, including Amazon. The results vary between the studies from a significant percentage of users concerned with privacy [146] to only 2% [82]. Using thematic analysis, Linden et al. [220] performed a comparative analysis of customer reviews of human and pet wearables and found that very few privacy concerns were expressed about these technologies. However, apart from a few studies for wearables, we still observe a lack of quantification of S&P issues expressed in reviews of popular consumer IoT devices and also on the character of the S&P signals embedded within the reviews.

#### MANUFACTURERS AND SELLERS

Manufacturers have significant power in determining the S&P posture of IoT devices. As the primary designers and producers, they decide which S&P features are implemented, how frequently updates are released, and how transparently S&P information is communicated to users. However, manufacturers lack sufficient incentives to address security features since doing so would increase production costs, reduce battery life and delay the time to market [21, 38]. Moreover, due to diverse nature of the IoT supply chain, ranging from chip manufacturers and firmware developers to platform integrators and end-device assemblers, there is a tendency to assume that somebody else in the chain might have addressed security concerns creating fragmented accountability and inconsistent S&P implementation [159].

Nonetheless, recent studies on manufacturers have shown promising results on their increased accountability, especially with regards to releasing firmware patches. Nakajima et al. [162], in their pilot study, found that five out of the six IoT vendors released patches on time showing the IoT device manufacturers do prioritise S&P. Pérez et al. [183] find similar results with a larger pool of 104 vendors and highlight that IoT-centric vendors release more patches on time than non-IoT-centric vendors. Further, they find no significant relationship between vendor size and patch availability suggesting that even smaller manufacturers devote resources towards S&P.

Sellers are the entities that manage the transactions on e-commerce platforms. On large e-commerce platforms, the same product is sold by multiple sellers who may compete on factors such as speed of delivery or pricing. One advantage of e-commerce platforms is that they level the playing field and provide higher visibility even to smaller sellers [59]. However, in the context of S&P transparency initiatives, the responsibility for providing such information may fall on sellers — many of whom, particularly smaller vendors, may lack the awareness, capacity, or resources to do so effectively.

## REGULATORS

Regulators play a crucial role in addressing market failures and correcting information asymmetries in the IoT ecosystem. As IoT devices become more pervasive and deeply embedded into different parts of our online and offline life, regulatory bodies across jurisdictions have introduced initiatives to improve and standardise their S&P features. Many of these initiatives aim to provide assurance – either directly from manufacturers or through third-party evaluation – that a device meets baseline S&P requirements. In addition they also promote transparency about S&P features to consumers enabling more informed purchase decisions.

Nation-level codes of practice, such as those introduced in the UK [217, 218] and the US [78, 150], represent foundational steps toward this goal. However, research highlights considerable challenges in such standardisation efforts, including the difficulty of achieving consensus across diverse stakeholders and of measuring the real-world effectiveness of such standards [121, 136]. More recent regulatory developments have focused on update support duration as a tangible and meaningful security commitment from manufacturers. The US Cyber Trust Mark, a voluntary labeling initiative expected to roll out in late 2025, includes update support duration as one of its key criteria [79]. Similarly, the UK's Product Security and Telecommunications Infrastructure (PSTI) Act mandates that manufacturers inform consumers about the duration for which updates are guaranteed for their IoT devices [66]. In contrast, the Cyber Resilience Act (CRA) in the European Union takes a stronger stance mandating that - in addition to update duration disclosure - security updates must be provided for a device's entire expected use period but at least for five years [74].

Other initiatives have placed the onus on online stores, which serve as the primary consumer interface during the purchasing process. These platforms, particularly larger ones, have more bargaining power than consumers to demand S&P information from manufacturers. In 2020, the Netherlands Authority for Consumers and Markets (ACM) began requiring Dutch online retailers to display update support duration information. This policy refers to EU directives on digital content and sales of goods [75, 76] as the legal basis.

The EU Energy Labelling Regulation provides a precedent for integrating S&P into existing labeling infrastructure. Under this regulation, large appliances including IoT devices like smart TVs are required to disclose update durations along with the energy information and submit them to the European Product Registry for Energy Labelling (EPREL), a publicly accessible database [3].

## E-COMMERCE PLATFORMS

The e-commerce platforms play a central role in shaping the environment in which consumers make their purchase decisions. These platforms are not passive intermediaries that merely connect buyers and sellers. Their design choices, from the user interface to the algorithmic recommendations, actively inform consumer decision making [43]. The platform algorithms that determine the product rankings and recommendations can amplify or suppress certain signals. For instance, a device reported as being hacked in reviews may still appear prominently in search results due to high sales volume or overall ratings. Moreover, the platforms determine the format and placement of product

information, such as technical specifications and product description. In addition, e-commerce platforms also play a gatekeeping role in enforcing (or neglecting) regulatory and industry standards. Without the platform support for provisioning, there would be no consistency even within the same platform on how regulatory compliance information is displayed which would make it harder for consumers to access the information at the time of purchase. However, despite e-commerce platforms being a critical focal point in reducing the information asymmetry there is little research that explores their roles in encouraging more S&P conscious purchase decisions.

## 1.5. RESEARCH GAPS

While prior research has significantly advanced our collective understanding on the technical, regulatory and user-centric aspects of IoT S&P, there is little security research on the e-commerce platforms where these purchase transactions occur. Studying these platforms and the signals embedded in the purchase context is important not only because the moment of purchase is an opportune moment of change [177], but also because these are not neutral platforms that merely connect consumers and sellers. They are active socio-technical arrangements that mediate consumer preferences and influence purchase decisions through sorting algorithms, product framing, review systems and the information presented and highlighted [43]. In the context of IoT purchase decisions, it is therefore crucial to understand how S&P-related information is communicated — or not — within these e-commerce platforms where most consumers now make their purchase decisions. This would allow us to identify where improvements are needed to facilitate more S&P-aware IoT purchases, because these platforms might not be naturally efficient at addressing the S&P preferences of consumers, they require careful design to ensure effective matching between consumers' S&P preferences and their IoT device purchases [191]. In this section, we outline the various gaps identified in the literature around S&P and the actors in the e-commerce platforms.

First, existing research has largely overlooked e-commerce platforms as a potential source of organic S&P signals. For example, the product description of IoT devices on their product pages might include (or not) crucial S&P signals like update support duration, a concrete and policy-relevant measure of manufacturer commitment to S&P. Consumer reviews might contain organic comments on the S&P aspects of IoT devices reflecting broader market sentiment. These organically occurring signals could provide a more accurate reflection of real-world consumer awareness and concern than experimental studies or controlled surveys.

Next, although studies have identified that consumers are interested in and willing to pay more for IoT devices with better S&P, it remains unknown if these preferences are reflected in their real-world purchase decisions especially in the context of the information asymmetry in the market. This limits our understanding of whether consumers reward manufacturers of IoT devices with better S&P, a factor that could potentially shift manufacturer incentives.

From a manufacturer's perspective, it is reasonable to expect that significant advancements in S&P features would be leveraged as a competitive advantage and clearly communicated to consumers. One possible explanation for the lack of information on S&P features on e-commerce platforms could be a lack of such significant advancements

over time. However, we do not know if this is indeed the case since there is a notable lack of longitudinal, empirical research examining how S&P features in IoT devices have evolved over time. The lack of such longitudinal data limits our understanding of how manufacturers are responding to security incidents with IoT devices or responding to the increased policy push for better S&P.

Finally, the underlying reason to focus on improving S&P of IoT devices is the severe consequences, especially societal consequences due to DDoS attacks from IoT botnets. However, DDoS attacks have existed long before IoT botnets and as outlined in [section 1.2](#), it remains unclear who suffers the negative externalities from IoT botnets or how the victim landscape has changed due to the advent of IoT botnets. This understanding is crucial because a change in the victimisation patterns might trigger a corresponding change in incentives.

In sum, significant gaps remain in our understanding on the presence of S&P signals in the market for IoT devices. From the consumer perspective, it is unclear to what extent concerns, comments and preferences about IoT S&P are organically expressed in reviews and also on whether S&P preferences of consumers – identified through user studies – translate to real world purchase decisions. From the perspective of manufacturers and sellers, there is limited data on how crucial S&P related signals like update duration is disclosed on the online stores. There is also no empirical study on whether the S&P of IoT devices has improved or remained stagnant over time. Furthermore, it is not yet clear who suffers the negative externalities from large-scale IoT-based DDoS attacks. These gaps highlight the need for a systematic analysis on the signals for IoT S&P and also on the dynamics of consumers, manufacturers within an evolving threat and regulatory landscape.

## 1.6. RESEARCH AIMS AND QUESTION

Understanding market signals for IoT S&P is crucial because they represent the lived reality of how consumers currently make purchase decisions in the absence of comprehensive regulation. While laboratory studies can measure willingness to pay under controlled conditions, and regulatory frameworks can mandate disclosure requirements, neither approach captures how information actually flows – or fails to flow – in real-world e-commerce environments where millions of consumers make daily purchase decisions. For regulators and policymakers, market signals provide evidence about where interventions are most needed: Are consumers already expressing S&P concerns but lacking reliable information? Are manufacturers improving S&P over time but failing to communicate these improvements? Do regulatory disclosure requirements actually reach consumers, or are they lost in the noise of e-commerce platforms? Most importantly, market signals offer a foundation for designing interventions that work with, rather than against, existing consumer decision-making processes and market dynamics.

This perspective – examining IoT S&P through the lens of market signals on e-commerce platforms – represents a fundamentally new angle in security research. While prior work has separately studied technical vulnerabilities, user perceptions, and regulatory frameworks, this dissertation connects these threads by examining how S&P information is actually communicated, perceived, and acted upon in real markets. This approach reveals gaps and opportunities that are invisible when studying any single actor in isolation:

the disconnect between manufacturer claims and retailer disclosures, the divergence between stated willingness to pay and actual purchase behaviour, and the variation in how different platform designs surface or suppress S&P information. By bringing the market perspective to the foreground, this research illuminates a missing piece in our understanding of IoT S&P – not just what security features devices have or what consumers say they want, but how information about S&P flows through the complex socio-technical system of e-commerce and shapes real purchase outcomes.

The aim of this dissertation is to conduct a comprehensive, empirical investigation into the S&P signals in e-commerce markets that sell IoT devices, focusing specifically on consumer-facing IoT devices – such as smart cameras, printers, speakers, TVs, and wearables – that individuals purchase for personal use in domestic settings. We deliberately exclude business-to-business (B2B) IoT deployments, industrial IoT applications, and connected vehicles, as these operate under different procurement processes, risk profiles, and regulatory frameworks. Our focus on consumer IoT is motivated by several factors: these devices are purchased by individuals who typically lack technical expertise to assess security features; they are sold through retail channels where information asymmetries are most pronounced; and they represent the fastest-growing segment of the IoT market where individual purchase decisions aggregate into significant market signals. Moreover, consumer IoT devices have been at the centre of high-profile security incidents – from privacy breaches to large-scale botnet attacks – making them both socially relevant and policy-critical to study.

This study seeks to identify both explicit and implicit S&P signals as they appear across the retail environments — including product descriptions, technical specifications, update support disclosures, and user-generated content such as customer reviews. By analysing these signals at scale, we will assess the extent to which S&P related information is visible and accessible to consumers at the point of purchase. In addition, the dissertation will explore whether and how consumer preferences for S&P are reflected in their real life purchase actions. It also includes a longitudinal component that evaluates the signals from IoT manufacturers in terms of changes in S&P features over time. Finally, it also presents an analysis on the changes in victimisation pattern of DDoS attacks from IoT botnets to identify the players who are most affected and therefore with highest incentive to improve the S&P of IoT devices.

Through this multi-dimensional analysis, this dissertation aims to fill a critical gap in our current understanding by linking regulatory and consumer facing aspects of IoT device security within the real world context of e-commerce platforms. The findings can help policy makers and stakeholders understand the pulse of the market with respect to IoT S&P, and understand how to better design our e-commerce platforms to facilitate more S&P-aware IoT purchase decisions. Overall, this dissertation aims to answer the following overarching research question.

*What signals for security and privacy are present in the e-commerce platforms that sell IoT devices?*

This research question is broken into several manageable sub research questions that have been addressed through five studies.

### 1.6.1. STUDY 1 – S&P SIGNALS WITHIN CONSUMER REVIEWS

Despite growing evidence that consumers care about secure IoT devices, relevant security and privacy-related information is unavailable at the point of purchase. While initiatives such as security labels create new avenues to signal a device's S&P posture, we analyse an existing avenue for such market signals - customer reviews. We investigate whether and to what extent customer reviews of IoT devices with well-known S&P issues reflect these concerns. This study addresses the research question: **(RQ1)** To what extent do customer reviews of IoT devices reflect concerns about security and privacy, and what specific themes emerge when such concerns are articulated?

We examine 83,686 reviews of four IoT device types commonly infected with Mirai across all Amazon websites in English. We perform topic modelling to group the reviews and conduct manual coding to understand (i) the prevalence of security and privacy issues and (ii) the themes that these issues articulate. We find that around one in ten reviews mention S&P, and distil these references into seven themes and two orthogonal themes. We draw on our results to make recommendations for online market places to promote more S&P conscious purchase decisions and identify future research directions.

### 1.6.2. STUDY 2 – IOT DEVICE SALES AND S&P SIGNALS

Although recent studies show that consumers care about IoT S&P, there are no empirical studies examining whether these preferences for IoT S&P translate into real world purchase decisions. We study the relationship between the S&P of IoT devices and their sales, considering the S&P signals in the context of these sales. We obtained expert S&P ratings of IoT devices from a European consumer association and the corresponding sales data from a leading Dutch online store. We complemented this with additional information like user ratings, the number of reviews and update support duration from two Dutch online stores.

This study answers the following research question: **(RQ2)** To what extent do expert evaluations of IoT device security and privacy align with consumer-facing signals — such as user ratings, sales performance, and the availability of update support information — on e-commerce platforms? Our regression model shows that a one-standard-deviation increase in S&P ratings corresponds to a 56% sales boost, and a correspondence between price of the device and its sales. Additionally, presence of update support duration corresponds to higher S&P ratings and is linked to a 69% increase in sales, suggesting positive incentives for affordable IoT devices with clear S&P information.

### 1.6.3. STUDY 3 – THE EVOLUTION OF IOT S&P

Inspite of significant efforts to enhance the S&P of IoT devices in recent years, concerns about their S&P remain critical. A compelling question persists: are newer IoT devices genuinely better equipped with S&P features than their predecessors? In this study, we address this gap by leveraging longitudinal S&P testing and rating data from a European Consumer Association known for conducting rigorous, expert-driven assessments of IoT S&P. We focus on three widely used and mature IoT device types – IP cameras, smart printers and smart speakers – and analyse their S&P trends over the past decade.

This study addresses the research question: **(RQ3)** How have the S&P ratings and features of three types of IoT devices evolved over the past years? Our analysis shows that

S&P ratings have remained broadly stable over the year. We also uncover a surprising inconsistency in the deployment of S&P features in successive IoT device models from the same manufacturer. We reflect on the possible causes and make recommendations for manufacturers to systematically improve the S&P of their IoT devices.

#### 1.6.4. STUDY 4 – IoT UPDATE INFORMATION ON ONLINE STORES

Even though there is legislation in several countries requiring increased transparency to consumers on the duration of update support for IoT devices, there is limited empirical research on what information is actually available to consumers. To address this gap, this study addresses the research question: **(RQ4)** To what extent is update support duration information available on EU online stores?

We conduct a large-scale empirical study on the disclosure of update support duration via product pages on online stores in the EU market: popular regional stores, EU Amazon stores and the Chinese e-commerce platform Temu. We conduct manual and automated analysis on 4600 and 22,298 product pages respectively for five popular IoT device types: IP cameras, smart printers, smart speakers, smart TVs, and smart watches. For a smaller sample, we compare the available information on online stores to the manufacturer websites. For smart TVs, we also compare it to a centralized EU product database where sellers are required to register energy efficiency and update support details before sale. Our results shows that stated support durations vary across retailers, manufacturers, and the EU's central product database. We find widespread contradictions, with retailers often understating support relative to manufacturers.

#### 1.6.5. STUDY 5 – VICTIMS OF IoT BOTNET ATTACKS

There is limited data on who suffers the negative externalities from IoT botnets, or in other words, the victims of IoT botnet powered DDoS attacks. The advent of IoT botnets like Mirai has challenged the popularity of older techniques like amplification attacks. In this study, we characterise the consequences of this change for the victimisation pattern of DDoS attacks. We conduct the first empirical comparison of victims of amplification attacks and botnet attacks and draw on the properties of targets outlined in Routine Activity Theory (RAT) to characterise the differences. We analyse the differences in the victimisation patterns at the level of IP addresses and Autonomous Systems. We further use RAT to outline the policy implications of our analysis.

This study answers the research question: **(RQ5)** Who are the victims of DDoS attacks using IoT botnets, and how does the victimisation pattern of IoT botnets compare to that of earlier attack techniques? The study shows that there are differences in the victimisation patterns at the level of IP addresses and Autonomous Systems with botnet attacks targeting more high profile victims.

**Table 1.1:** Dissertation outline.

Chapter	Publication
Ch. 2	<b>Vetrivel, S.,</b> Van Harten, V., Gañán, C. H., Van Eeten, M., & Parkin, S. (2023). Examining Consumer Reviews to Understand Security and Privacy issues in the Market of Smart Home Devices. In <i>Proceedings of the 32nd USENIX Security Symposium (USENIX security 23)</i> (pp. 1523-1540).
Ch. 3	<b>Vetrivel, S.,</b> Bouwmeester, B., van Eeten, M., & Gañán, C. H. (2024). IoT Market Dynamics: An Analysis of Device Sales, Security and Privacy Signals, and their Interactions. In <i>Proceedings of the 33rd USENIX Security Symposium (USENIX Security 24)</i> (pp. 7031-7048).
Ch. 4	<b>Vetrivel, S.,</b> van Eeten, M., & Gañán, C. H. Time Will Tell: A Longitudinal Analysis of the Evolution of Security and Privacy of Three IoT Device Types. ( <i>Under Submission</i> )
Ch. 5	<b>Vetrivel, S.,</b> van Eeten, M., & Gañán, C. H. Missing, Present and Conflict-ing: A Large Scale Analysis of IoT Update Information in the EU Market. ( <i>Under Submission</i> )
Ch. 6	<b>Vetrivel, S.,</b> Noroozian, A., Makita, D., Yoshioka, K., van Eeten, M., & Gañán, C. H. (2023). Birds of a Feather? A Comparative Analysis of DDoS Victimization by IoT Botnet and Amplification Attacks. In <i>22nd Workshop on the Economics of Information Security (WEIS 2023)</i> .

# 2

## ANALYSIS OF CONSUMER REVIEWS

*Despite growing evidence that consumers care about secure Internet-of-Things (IoT) devices, relevant security and privacy-related information is unavailable at the point of purchase. While initiatives such as security labels create new avenues to signal a device's security and privacy posture, we analyse an existing avenue for such market signals - customer reviews. We investigate whether and to what extent customer reviews of IoT devices with well-known security and privacy issues reflect these concerns. We examine 83,686 reviews of four IoT device types commonly infected with Mirai across all Amazon websites in English. We perform topic modelling to group the reviews and conduct manual coding to understand (i) the prevalence of security and privacy issues and (ii) the themes that these issues articulate. Overall, around one in ten reviews (9.8%) mentions security and privacy issues; the geographical distribution varies across the six countries. We distil references to security and privacy into seven themes and identify two orthogonal themes: reviews written in technical language and those that mention friction with security steps. Our results thus highlight the value of the already existing avenue of customer reviews. We draw on these results to make recommendations and identify future research directions.*

### 2.1. INTRODUCTION

Among the range of consumer IoT devices now available – like smart doorbells and smart home surveillance systems – many lack sufficient security and privacy features that are fit for purpose. These shortcomings have been exploited in various ways, most visibly in large-scale Distributed Denial of Service (DDoS) attacks from compromised IoT devices [23, 124]. These emerge as part of a broader trend to leverage vulnerable IoT devices for malicious activities, ranging from botnets as criminal infrastructure [213] to cryptojacking [32], and intimate domestic abuse [36, 211].

A broad consensus has emerged that the root cause of this trend is a market failure caused by incentive misalignment [196]. Manufacturers of the affected devices do not have sufficient incentive to improve security. The cost of poor security is not borne by them but by others, such as the network operators and service providers that suffer

DDoS attacks. However, the incentives for manufacturers would improve if customers care about and prioritise security and privacy in their IoT device purchases - and recent research demonstrates that users not only care about IoT security and privacy but are also willing to pay for it [34, 89]. The failure then is one of information asymmetry, a 'market for lemons' [14] where the consumer cannot discern good from bad with the information available.

In order to fix this market failure and hasten improvements to the security and privacy ('S&P' from here onwards) of IoT devices, external interventions have been pursued. To that end, there are various efforts to introduce a range of *market signals* [191, 204] - more information about the quality of the device - to reduce the *information asymmetry* between S&P observers (such as experts, governments) and consumers. These signals can be marketing-controlled, within the control of the manufacturer like security 'labels', or non-marketing-controlled like customer reviews [13].

Self-certified security labels that state assurances [69, 72, 157] and aim to standardise the S&P information available to consumers - like food nutrition labelling - are still in formative stages and are not yet in use. Further, it is not assured that they would reflect the actual real-world security posture of the device or directly respond to the concerns that consumers have [86]. Crucially, these efforts overlook how much the consumer base is already recognising - and signalling - a need for S&P in IoT devices and what the expression of those needs looks like. Such signals are present within non-marketing-controlled avenues like customer reviews.

Our primary motivation in this study is to evaluate if customer reviews - as an existing source of signals about the quality of consumer IoT products - voice consumers' S&P-related concerns. Given that prospective buyers look at reviews when considering a purchase [240], understanding the content of reviews can inform the views that consumers express and also how accessible those views are to other consumers as signals of device quality. In turn, this can point to where action is needed to further inform decisions around the purchase of a smart device.

Moreover, if S&P issues raised in customer reviews can be characterised and amplified, they will highlight the S&P posture of IoT devices and potentially incentivise manufacturers through the impact on brand reputation [27]. These signalling effects are agnostic to whether the S&P-issues in the reviews are trustworthy, misinformed, or even whether the review is fake [97], as long as prospective buyers trust the reviews as an important source of information which research suggests they do [240].

Here, we present a large-scale investigation and characterisation of signals for S&P, through an analysis of consumer reviews, in the market of IoT devices commonly infected with 'Mirai-like' malware [23, 189] and related device features. This device selection allowed us to see to what extent S&P-issues notorious within the security community have permeated the marketplace and are raised - unprompted - by consumers.

We address two research questions: **(RQ1)** What fraction of customer reviews for IoT devices articulate security or privacy issues? **(RQ2)** When security issues or privacy are mentioned, what themes are being articulated? To answer these questions, via distinct sampling strategies, topic modelling and qualitative analysis approaches, we collected and analysed 83,686 reviews from all of the six country websites of Amazon, which are natively in English. The scraping process from product link collection, cleanup and

subsequent review data collection was conducted between May and August 2021.

Our main contributions are as follows:

- We present a comprehensive evaluation of security and privacy (S&P) ‘excerpts’ for consumer IoT devices, as represented in consumer reviews. Consumer reviews are an intrinsic part of purchase deliberation, and as such, form a critical intervention point for informing improved security purchase decisions. Further, these reviews are the unprompted S&P views of consumers within a mix that includes other non-S&P preferences.
- We show that approximately 10% of the reviews contain security-related issues, which means that prospective buyers have a limited chance of encountering this information when considering various products for purchase.
- We distil the themes in reviews that articulate security and privacy concerns. We find that S&P information is articulated both in technical and non-technical terms and spans a variety of themes, from firmware updates to worries about data capitalism.
- We describe a novel combination of machine learning to categorise IoT product reviews and manual thematic analysis of the text to understand context, pain points, and themes in reviewers’ own terms.
- We discuss several recommendations for strengthening the information and signalling value of customer reviews and leveraging this existing mechanism to reduce information asymmetry and the security incentives of manufacturers.

## 2.2. BACKGROUND AND RELATED WORK

Here we provide a background on online reviews and describe the existing research on user perceptions of IoT S&P and how this fits into a consumer/market context.

### 2.2.1. ANALYSIS OF ONLINE REVIEWS

Online product reviews have been recognised as containing critical information regarding consumers’ concerns [134] and are considered essential in building a firm’s business intelligence [236]. Reviews have been noted to function both as informants and as recommenders influencing both product sales and purchase decisions [176]. Moreover, studies on the impact of online reviews show that the quality and quantity of product reviews positively influence the sales of a product [139]. With respect to reviewers themselves, Hu et al. [102] show the self-selection bias at play where buyers with extreme positive or negative experiences are more likely to post online reviews. However, Han et al. [93] report that this bias is mitigated when buyers are familiar with the online review platform. Given the popularity of Amazon, we expect that most buyers are familiar with its review platform and are, therefore, more likely to post reviews for varying levels of satisfaction. This could also explain why Amazon is the most popular source of customer reviews within the research community [216].

Moreover, customer reviews of other devices, including certain types of IoT devices like smart home assistants and wearables, have been studied previously. Two such studies [82, 146] use unsupervised machine learning techniques to understand if users express

any privacy concerns in reviews of popular e-commerce sites, including Amazon. The results vary between the studies from a significant percentage of users concerned with privacy [146] to only 2% [82]. Using thematic analysis, Linden et al. [220] performed a comparative analysis of customer reviews, also on Amazon, of human and pet wearables and found that very few privacy concerns were expressed about these technologies. However, in our work, we go beyond merely reporting S&P concerns, quantifying the presence of S&P issues in reviews and also qualitatively determining the character of S&P signals within these reviews.

### 2.2.2. CONSUMER PERCEPTIONS OF IOT SECURITY AND PRIVACY

Earlier work on consumer perceptions of IoT S&P has primarily been conducted through surveys, semi-structured interviews and experience sampling. To examine the mental models of users of smart devices, Abdi et al. [7] and Zeng et al. [235] conducted semi-structured interviews and reported gaps in users' mental models regarding security. They attribute these gaps primarily to limited technical understanding and point to ad-hoc (and typically non-technical) strategies employed by users in order to protect themselves.

With respect to privacy, a study by Williams et al. [229] observes that the price of consumer IoT devices deterred more users than privacy concerns and note that since the purchase of IoT devices is voluntary, privacy was more likely to be sacrificed for functionality rather than necessity. This trade-off is echoed in other studies [96, 239], which observe that users balance the risks of using IoT devices against the convenience and benefits offered. In our analysis of reviews, we find issues which essentially revolve around consumers 'not knowing what they were getting into' when purchasing a device, uncovering challenges in using and understanding their new device and how it fits into their smart home environment.

### 2.2.3. INTERVENTIONS IN CONSUMER IOT PURCHASE DECISIONS

Prior interviews with retail customers have highlighted the point of purchase as a critical point for informing decisions about the security of new computing devices [177]. In making decisions about security, home users get their security and privacy advice from various sources, where this can include family, friends, and peers [63, 186], informal technical experts [172, 180], and media such as news stories, blogs, and TV [64, 185].

Studies have analysed the marketplace for IoT devices to explore methods and mediums to inform consumers of the privacy posture of devices and their corresponding consequences. For instance, Gopavaram et al. [90] investigated customers' Willingness-To-Pay (WTP) for privacy vs their Willingness-To-Accept (WTA) a lack thereof through an emulated marketplace study. WTA participants, presented with the highest privacy settings by default, were more likely to pay a premium and purchase devices with a higher privacy rating, thereby indicating that interface design also influences purchase decisions.

Along similar lines, Blythe, Johnson and Manning [34] note that providing people with simple security-related information prior to making their purchase decision has the potential to encourage the purchase of devices which are more secure. They also found that consumers are willing to pay more for increased security and that the relative amount of risk reduction has no significant impact on that willingness.

In addition, various forms of S&P information labels have been proposed, includ-

ing a graded label, labels indicating S&P features, and labels indicating ‘approval’ by independent assessment [111]. Blythe et al. [111] found that except for cases where an information label indicated that a device has poor security, consumers were significantly more likely to buy a device with a label. They also found that although functionality was generally more valued, people were willing to pay the same premium for both improved functionality and security. Emami-Naeini et al. [72] presented participants with labels which were a mix of the aforementioned label types and also reported positive feedback from participants who indicated that they struggled to find this kind of information at the time of device purchase.

With a focus on the availability and regularity of security updates for IoT devices, Morgner et al. [157] found support among their survey participants for this kind of information, more so for those who perceived higher risks in using such devices. Broader efforts to both improve and standardise IoT S&P features include nation-level codes of practice (including in the UK [217, 218] and US [78, 150]), where various challenges to such device standardisation efforts have been highlighted in research [121, 136], including agreement on standards and evidencing their effectiveness. Many of these interventions seek to standardise various assurances, either from manufacturers themselves or from independent experts, that the S&P properties of a device are sufficient to be able to use a newly-purchased home IoT device securely.

Here we explore the signals and indicators of S&P issues which emerge from owners themselves, as expressed in reviews in the setting of an online shopping platform. These not only identify concerns but also ‘hotspots’ in device use where these issues become critical and S&P expectations which were not met. Moreover, our findings relate issues of awareness and preferences around S&P to the availability of information for an adequately informed purchase.

## 2.3. METHODOLOGY

This section outlines our data collection and analytical approach. Our starting point for collecting the reviews is the online marketplace and shopping website Amazon. Amazon is a dominant e-commerce platform with a large customer base across different countries.

The libraries used for topic modelling work better in English, so we only sourced customer review data from Amazon websites that are natively in English: amazon.com (United States), amazon.com.au (Australia), amazon.ca (Canada), amazon.in (India), amazon.sg (Singapore), and amazon.co.uk (United Kingdom).

### 2.3.1. SELECTION OF IOT DEVICES

Research on IoT malware, botnets and compromised devices has identified specific products being compromised at scale because of severe security failures, such as using known factory-default credentials. Many of these devices fall into four categories [189]: surveillance systems (including DVR/NVR), set-top boxes, smart home hubs, and routers. Although routers are often not considered as being an IoT device, they are integral to home networks and also susceptible to IoT-related attacks.

We approached device selection in two ways. First, we searched for specific products known to be vulnerable to IoT malware infections (specifically Mirai) [189]. This produces

a set of devices with a high likelihood of prior, if not still current, security issues. We were interested to see if the reviews for these products would contain more comments on security or privacy (S&P) than other products for the same device type. We searched for these once-vulnerable devices on the Amazon websites using a combination of manufacturer, model name/number and device type (e.g., XXX YYY-123 Router, device list from Appendix A of [189]) and collected the product links. Of the 53 IoT devices we searched for, only 16 were still being sold: 14 routers, one DVR, one set-top box, and no smart home hub. The DVR and set-top box did not contain any reviews and were dropped, resulting in a single category of 14 once-vulnerable routers.

As the second part of device selection, we expanded our search to additional products in the same four product types commonly infected with Mirai (surveillance systems, set-top boxes, smart home hubs, and routers). Since these products do not fall under a single category on Amazon, we used different search terms. We searched for a set of commonly used terms for each device type. For instance, surveillance systems might be referred to as surveillance cameras, IP cameras, security cameras, etc. The terms used to search for each device type are added to the Appendix A.1.

### 2.3.2. PRODUCT PAGE AND REVIEW RETRIEVAL

A Python script was written to search for the terms detailed above and collect all the product links on the first page of search results, along with the partial product title visible on the page to help with the manual cleanup in the next step. We repeated this for each of the six country websites.

The script returned a total of 3,524 links across all six websites and four device types. The number of product links differs per category because we used more queries for some categories than others (as shown in Table A.1 in the Appendix). For example, for surveillance systems, we also included DVRs and NVRs and thus included queries for those. However, the different numbers of product links per category has no impact on our analysis, as we analyze each category separately to answer our research questions.

Before scraping the reviews for these products, we first reviewed all product links manually. More than 60% of them were dropped because they did not point to an IoT product. For instance, across IP cameras, there were multiple results for fake cameras that merely act as a deterrent for burglars and cameras that do not connect to the internet. With a focus on internet-enabled devices, we excluded products that do not connect to the internet or only use their proprietary mesh network for connectivity. Devices with optional internet connectivity, like IP cameras that could be connected to the internet using a sim card, were kept in the set. Likewise, results for devices like baby monitors, spy cams, and pet cams that were internet enabled were retained. The final set consisted of 1415 product links. The count of product links collected per Amazon website and device type is Table 2.1.

### 2.3.3. REVIEW DATASET CONSTRUCTION

Using the product links we collected, we scraped the customer reviews for each product using Python scripts. For each product link, we collected two sets of reviews—one set for each research question (Figure 2.1). Our first research question is ‘What fraction of customer reviews for IoT devices articulate security or privacy issues?’. To answer this,

**Table 2.1:** Count of collected product links for each device type after cleanup.

Amazon website country	Surveillance systems	Routers	Set-top boxes	Smart home hubs	Once-vulnerable routers	Total
Australia	130	72	1	12	5	220
Canada	271	123	24	5	22	445
India	87	45	1	2	1	136
Singapore	76	42		5	3	124
UK	185	40	20	18	10	265
USA	132	50	8	17	9	207
Total	881	372	54	41	68	1415

**Table 2.2:** Count of reviews collected for each research question.

Amazon website country	Surveillance Systems		Routers		Set-top boxes		Smart home hubs		Once-vulnerable routers		Total	
	RQ1	RQ2	RQ1	RQ2	RQ1	RQ2	RQ1	RQ2	RQ1	RQ2	RQ1	RQ2
Australia	940	147	886	87	0	0	667	94	0	0	2493	328
Canada	12102	1452	8198	799	679	61	120	18	442	58	21541	2388
India	4189	630	4900	720	5	2	145	16	67	2	9306	1370
Singapore	38	8	113	11	0	0	0	0	0	0	151	19
UK	11908	2206	3812	1356	1178	123	53	9	669	153	17620	3847
USA	10626	4207	4913	2462	319	36	506	113	1037	404	17401	7222
Total	39803	8650	22822	5435	2181	222	1491	250	2215	617	68512	15174

we scraped the first 30 reviews (three pages) from each of the five star ratings for each product link (from one star to five stars). This resulted in 150 reviews for most products, though for some products, there were less than 30 reviews with a particular star rating, leading to a slightly smaller set.

The upper limit of 30 reviews or three pages is in line with the results of a market research study [149] that showed that only 8% of consumers read more than 26 reviews before purchase. We chose this to reflect the number of reviews a realistic, motivated buyer would read before making a product purchase. Since other studies [81] show that, on average, seven reviews are read before purchase, we found this approach better suited to study what S&P-issues a prospective buyer might encounter than collecting all product reviews.

Moreover, we wanted to collect diverse reviews that remain agnostic to the actual skew of the rating distribution since popular reviews and ratings tend to have a self-reinforcing effect. Therefore, by design, we used a uniform sampling strategy and collected reviews across different star ratings in the order that Amazon displayed them, which is also the order in which a prospective buyer would see them. This was done to avoid bias in the review collection and to include reviews associated with different customer experiences. However, despite collecting reviews over all the star ratings, the final results showed a slight skew towards five-star ratings.

Our second research question is ‘When security or privacy issues are mentioned, what themes are being articulated?’. For this question, we wanted to focus on reviews that explicitly mention security and privacy issues. Since the top reviews collected to answer

the first research question are not representative of S&P-issues raised in reviews, we used Amazon's 'search customer reviews' option for each product link. The keywords used for searching were informed by prior user studies [7, 87, 96, 229, 235, 239]. This was done to account for how users articulate these concerns rather than the 'tech-savvy' words used within the research community (e.g., a user may say 'setup' instead of 'configuration'). Even though terms like "get into", "always listening" and "big data" were used by users in those studies, we did not include them in our queries as Amazon does not support concatenated search terms. The list of keywords is presented in Appendix A.2, grouped by category. For all scraped reviews, we collected the title, content, date the review was posted, the country it was posted from and the number of people that voted the review helpful. The username was not collected because of privacy considerations. The number of reviews collected for each research question across each website and device type is shown in Table 2.2.

#### 2.3.4. QUANTIFYING PRESENCE OF IOT S&P ISSUES

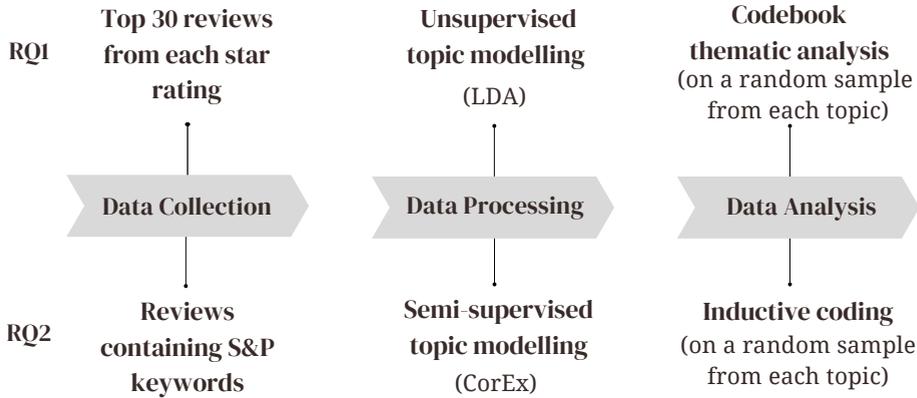
To answer our first research question – regarding the extent to which S&P issues were being discussed in reviews – we take a two-step approach. We first conduct unsupervised topic modelling on the reviews in each device type category to arrive at coherent clusters of reviews around specific topics. We then draw a random sample of 50 reviews from each topic and manually classify each review as to whether it discusses security or privacy or not. This provides us with a quantification of what portion of the reviews for a particular device type mention security or privacy issues.

We used Latent Dirichlet Allocation (LDA) topic modelling to discover 'topics' in the review dataset [33]. Two user-defined parameters are entered into the LDA model - the number of topics ( $k$ ) and the number of words for each topic. Once a set of topics has been generated, a coherence test can be used to assess the quality of the results based on the distance between words in the same topic. However, since it is difficult to assess 'k' a priori, LDA was run for different 'k' values, and the value with the best coherence score was chosen for each model. In total, the models identified 22 topics—12 for Surveillance Systems, six for routers and two each for the remaining device types.

Next, to quantify the number of reviews related to security or privacy in each product category, we randomly sampled 50 reviews from each review topic, which amounted to 1100 reviews. By drawing the manual samples from the topics rather than from the complete set of reviews, we avoid specific, more prominent topics from dominating the random sample. This way, we get a better sense of the diversity of reviews. In addition, we drew 100 random reviews from each country to check the geographical distribution of S&P reviews, which amounted to 600 reviews. Since our intent here was to check for variation across countries and not per device, we collected random samples from a pool of all reviews from each country.

Two researchers manually labelled these 1700 reviews according to whether they discussed security or privacy-related issues or not. This activity followed a thematic analysis approach [50]. More precisely, we used 'codebook' thematic analysis to manually categorise the content as S&P-related or not; and inter-rater reliability does not impact the quality of these results [40, 148]. However, disagreements (7.2% for reviews from each topic and 6% for geographical comparison) in the classification were resolved through

discussion and clarification of what was within the scope of S&P. In addition to obvious statements about the security or privacy of a product, reviews which referred to, e.g., firmware updates and those with mentions of authentication during setup, were also considered within scope.



**Figure 2.1:** Overview of steps followed for each research question.

### 2.3.5. DETERMINING CONTEXT OF IOT S&P THEMES

After quantifying the presence of security and privacy (S&P) issues, we examined the themes articulated in these reviews. For this purpose, we collected 15,174 reviews via search queries with security and privacy-related keywords. We then ran the topic modelling technique to identify clusters from which we could sample reviews for qualitative thematic analysis. For this, we chose semi-supervised Anchored Correlation Explanation (CorEx) topic modelling [85]. Unlike LDA, the CorEx topic model makes few assumptions about the latent structure of the data and flexibly incorporates domain knowledge through the anchor words that are fed to the model.

The highest coherence value was obtained for a ‘k’ value of 8, and we, therefore, ran the Anchored CorEx for eight topics. Since we had six categories of search words, we anchored each topic to two search words from each category that had the highest number of search results. This was done to guide the model towards these search groups, while also leaving room for other related words to be picked up as part of the same topic. This allowed us to group the reviews into eight clusters based on the topics.

In the next step, a random sample of 100 reviews from each of these topic clusters was taken for thematic analysis. This sampling technique ensured that the samples taken for the thematic analysis were representative of each cluster. The thematic analysis allowed us to better understand the contexts in which S&P feature in customer reviews. The methodology outlined by Braun and Clarke [50] was followed for the thematic analysis

for the context of S&P issues. One coder analysed the produced dataset, performing inductive coding to identify themes emerging from the review content [40]. Regular review meetings were held with other researchers in the team to clarify and discuss the codes, which helped trim, extend or change codes as needed. Figure 2.1 provides an overview of the steps followed for each research question, from data collection to analysis.

The Atlas.ti qualitative data analysis software was used to code certain portions of text. The portions relevant to security and privacy in each review were assigned a code based on what they represented. These codes are included in Appendix A.5. Once this was done for all 800 reviews, in an iterative process, we grouped the codes into groups based on the underlying theme. These themes and the corresponding review counts are shown in Table 2.5.

### 2.3.6. RESEARCH ETHICS

The Ethics Review Board of our institution approved the study design and data management protocol. We evaluated our research design against the principles of the Menlo Report for ethical practices in computing studies [119]. Data was collected on publicly available customer ratings and reviews on Amazon websites, and the associated usernames were not collected. Moreover, the scraping process was distributed over a longer duration through added delays in the script to ease the load on Amazon servers. Further, the scripts were fed specific pages and were not crawlers. With respect to ‘justice’, our study design aims to contribute to reducing asymmetry for all consumers, not specific groups.

## 2.4. RESULTS

Here we present the results of our combined topic modelling and thematic analysis activities. Reviewers are indicated by R###, and a device type classification is indicated by an additional letter: Set-top Box (B), Router (R), Once-vulnerable Router (RO), Surveillance system (S), Home Hub (H). All quotes from reviews are included verbatim, including potential textual idiosyncrasies and errors.

### 2.4.1. S&P PREVALENCE IN REVIEWS

As outlined in Section 2.3.4, we manually analysed 50 random reviews from each of the 22 topics output by the LDA models to answer our first research question. An overview of these topics is added to Appendix A.3. We thus analysed 600 reviews for Surveillance Systems, 300 for routers, and 100 each for Hubs, Set-top boxes and Once-vulnerable Routers.

During the classification of the reviews, we encountered mostly straightforward references to the security or privacy properties of the devices. There were also borderline cases, such as customers mentioning the availability or lack of firmware updates to get new features for the device rather than for security purposes. This is where a ‘code book’-oriented approach to analysis was utilised, adjusting the definitions of what was classified. To avoid under-counting the presence of relevant security and privacy information in the reviews, we also classified borderline cases as containing relevant security and privacy information.

Another area of divergence was reviews referring to device setup. Many reviews comment on setup being either easy or difficult, e.g., “*I haven’t bought a router in a while to be honest but this was staggeringly easy to set up* (R20188-R)”. We only classified as security-related those reviews that refer to security steps during the setup process, like login, password and authentication. The reviewer might evaluate these security actions negatively: “*dvr forces you to put a password we don’t want passwords the software could have major revisions done to it to make easier to use* (R19314-S)”.

In the end, classification resulted in varying numbers of reviews per topic sample referring to S&P. The highest references were within the topics for routers (15/50). The lowest (5.9% for Surveillance Systems) is still higher than other results [82]; we revisit this discrepancy in the Discussion. On average, each sample had five reviews relating to S&P. This shows that S&P issues are not siloed in specific conversations and instead emerge in most contexts. We explore these contexts further in the next subsection.

The overall results of the classification for each device type category are shown in Table 2.3. Across all top reviews, about one in ten reviews (9.8%) articulate S&P-related issues. On the one hand, this is a minor fraction of all reviews. On the other hand, it does mean that potential buyers browsing reviews stand a decent chance of encountering comments on the S&P properties of the devices they are looking at. It also means that review writers feel these aspects are important enough to mention them one in ten times, which is a non-trivial amount given the brevity of most reviews: the average review in this dataset contains 100 words (median: 65). That is less than the length of this paragraph (123 words).

**Table 2.3:** Percentage of reviews referring to security and privacy issues for each device type.

Device type	% of S&P reviews
Surveillance systems	5.9
Routers	16.3
Hubs	8.7
Set-top boxes	6.4
Once vulnerable routers	13.6
Total	9.8

The percentage of S&P reviews varies across device types: only 5.9% of the reviews for surveillance systems, while for router reviews, it is almost three times larger (16.3%). These differences cannot be explained by significant thematic differences across devices. The themes discussed in the next section are present among all device types. It was found that routers trigger more reviewer comments on S&P than the other device types. This might reflect the awareness of reviewers of how crucial a router is to the overall security of the home network. Routers have also been targeted and compromised by IoT malware, but the same holds for surveillance systems. So it is not clear that the ongoing attacks explain the differences. The set of once-vulnerable routers, which have been compromised at scale, has a slightly lower prevalence of S&P (13.6%) than the general router category (16.3%). This difference is not statistically significant, suggesting that the known issues with these routers did not cause a substantial increase in security-related

comments in those reviews.

### S&P CONCERNS IN REVIEWS ACROSS GEOGRAPHIES

In order to compare S&P concerns across the six countries in our dataset, we drew a random sample of 100 reviews from each country and manually classified them as S&P related or not. In some cases, when there are not enough reviews for a product on the country's website, Amazon posts reviews from other countries. However, for this analysis, the reviews were chosen based on the country it was posted from rather than the Amazon website from which it was scraped. The results are presented in Table 2.4. The average across all the countries (9.5%) is in line with the overall results from the device-wise analysis (9.8%), and the standard deviation is 3.67%. The deviation can be explained by the results for the US, which is higher than average (16%) and India, which is half the average (5%). For Australia and Canada, the results are in line with the global average of 10%, while for Singapore and UK, the percentage is 2% less.

**Table 2.4:** Comparison of S&P concerns across the stores.

Amazon review country	% of S&P reviews
Australia	10
Canada	10
India	5
Singapore	8
UK	8
USA	16

### CORRELATION BETWEEN RATING ASSIGNED AND MENTION OF S&P

Over all of the 1700 reviews that we analysed manually (1100 from device-wise analysis and 600 from country-wise analysis), we checked for a correlation between the ratings given and the mention of the S&P issue. The results of a biserial-point correlation test show a weak negative correlation (-0.04), indicating that mention of the S&P issue is associated with a lower rating. However, the results were not statistically significant (p-value = 0.11). In addition, we observed that despite our uniform sampling strategy, the ratings were not normally distributed - there was a slight right skew towards higher, 5-star ratings.

#### 2.4.2. INDUCTIVE THEMATIC ANALYSIS RESULTS

To answer the second research question, 'What themes are being articulated in reviews with S&P comments?', 100 reviews were sampled randomly from each of the eight topics output by the semi-supervised algorithm (see Section 2.3.5). Thematic analysis was then conducted on these 800 reviews. The first step in the thematic analysis involved inductively defining and adding a single code to reviews based on their content. We defined 99 granular codes before reaching saturation. Of the 800 reviews we analysed, 485 (60.6%) did not contain any references to security or privacy, even though they contained one of our search terms. This was an expected side-effect of using organic terms of real

users, such as “record”, “log”, etc. This inevitably selects many reviews that are not related to security or privacy. These reviews did not receive a code and were not considered for further analysis. In the final step, we condensed the 99 codes into a smaller set of themes based on high-level commonalities among the codes. Appendix A.5 contains the full list of codes and themes.

After this analysis, two additional properties of the reviews stood out that were orthogonal to the substantive themes. First, some reviews were written in quite a technical language, referencing specific protocols or technical artefacts. For example, “... NOTE: If you value your privacy you should put these cameras in their own vlan with NO outgoing access to any other vlans or network...” (R640-S). In some cases, reviewers using technical language mention being in, or having experience with, IT. Other reviews try to explain issues without using technical terms, as in: “The website to activate this device was banned from my phone saying its not a secure site and potential threat. My phone and computer could not even enter their site due to risk.. Im returning asap it was useless.” (R952-R)

The second distinction which we observed was whether reviews expressed personal frustration and friction with the steps involved in security configurations, e.g., “It is app control. Every device [connect] need to open apps & new password setup which is so bother me.thanks”(R11708-R), or not, e.g., “...It just prompts you to scan a code that allows your phone to download the app. Then scan code again, you are up and watching your cameras on the phone...” (R11708-R). This distinction is about sentiment, as separate from whether the review has a positive or negative evaluation.

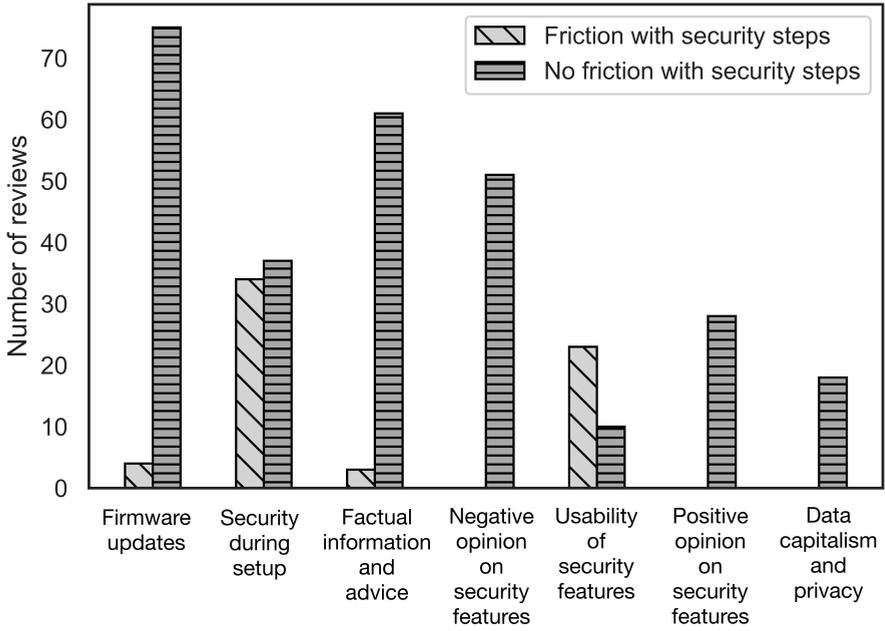
We added two additional codes to all reviews to complement the thematic code: whether they contained technical statements (yes or no) and whether they expressed friction with the security features (yes or no). The distribution of these two distinct themes over the seven themes is shown in Figures 2.2a and 2.2b. Of the 63 reviews (20%) that express friction with security steps, none were written in technical language. Although the technical reviews also mention experiencing trouble with security features, they tend to articulate problems as specific technical critiques of security features – one of our thematic labels. In contrast, reviews in non-technical language express the trouble as personal frustration for not being able to achieve the desired result.

Table 2.5 presents the distribution of reviews across the seven substantive themes and two distinctions of friction and technical nature. It also includes the average star rating of the reviews in each subset.

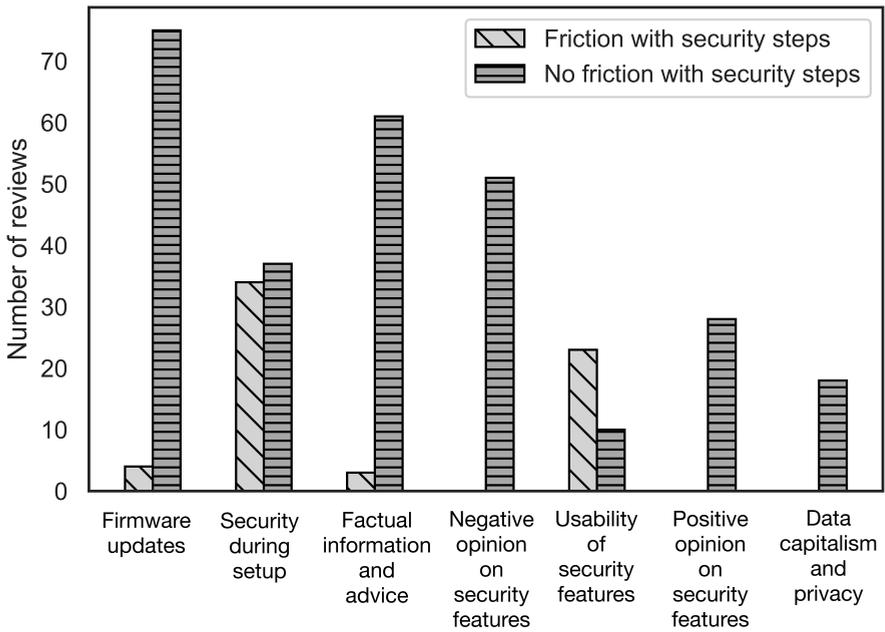
### FIRMWARE UPDATES

We begin the examination of the substantive themes with firmware updates. These updates are the primary medium for installing security patches on devices. Of 315 reviews with S&P issues, 78 refer to firmware updates. Only 14 of these talk about it positively. The rest are complaints about the update process.

As discussed in Section 4.1, we included all reviews discussing firmware, given its crucial role in device security and its potential relevance to prospective buyers. Of the 78 reviews in this theme, 23 talk about firmware updates in relation to issues with the device—complaints about updates not solving an issue, causing it or not helping with it. A couple of reviews mention being annoyed with frequent updates “too many interruptions for firmware patches” (R11035-R), while on the other hand, a handful are appreciative



(a) Distribution of technical and non-technical reviews over the seven substantive themes.



(b) Distribution of reviews mentioning friction and those not over the seven substantive themes.

**Figure 2.2:** Comparison of review distributions across different themes.

**Table 2.5:** Distribution of reviews over themes, friction, technical nature, and security-related versus not security-related, as well as average review score per subset.

Theme	Avg. rating	Count	%
Firmware updates	3.23	78	24.8
Security during setup	4.08	71	22.6
Factual information and advice	3.55	62	19.7
Negative opinion on security features	2.16	49	15.6
Usability of security features	2.68	33	10.5
Positive opinion of security features	4.75	27	8.6
Data capitalism and privacy	3.00	18	5.7
Friction with security steps	2.36	63	20.0
No friction with security steps	4.52	252	80.0
Technical reviews	3.52	102	33.7
Non-technical reviews	3.63	213	66.4
Security-related	3.41	315	39.4
Not security-related	3.56	485	60.6

of it, such as “*Pros: incredibly fast, frequent updates*” (R8927-B). There were also a few reviews about firmware being “*too buggy even after updating*” (R13798-R), firmware update process being “*unnecessarily kludgy*” (R11624-R) and firmware updates containing feature updates as well. Most of the reviews about issues during firmware update mention reaching out to customer service for assistance, but without always being able to resolve them:

*“...After going to the [manufacturer] website, I found that the solution was to update the firmware. [...] somehow during the update it got stuck or corrupted [...] I sat on the phone for over an hour with support...”* (R10576-R)

In assessing the usability of the update process itself for consumer IoT devices, Haney & Furman [94] found that participants experienced a lack of transparency in how updates worked and how. Where they noted a disconnect between updates and security, here we see a similar disconnect with firmware updates, specifically, around expectations of their role in resolving issues with devices.

Only eight reviews discuss firmware updates explicitly in relation to security. Some mention the auto-update feature as a security benefit (“*...Self updating. This system keeps itself up to date with the latest firmware, and software patches for stability and security...*” - (R13681-R) although one complained that “*...Autoupdate did not work...*” (R12665-R). Some reviews mention particular security vulnerabilities that they would like a patch for. However, only in one case was the patch available:

*“Was looking for a cheap router with updated Firmware available that included KRACK Patch ... I had to manually download and update the firmware...”* (R11045-R)

A similar desire for timely updates has been reported elsewhere in an exploration of user information expectations for product ‘security labels’ [157].

### SECURITY DURING SETUP

Device setup is also a phase of key importance in the security of purchased devices. 21.5% (172) of the 800 reviews refer to setup. Interestingly, most reviews that refer to setup express polar opinions on the spectrum of it being very easy to frustratingly complicated. This is more likely a reflection of user expectations regarding the setup process than a direct indication of its difficulty, i.e., people experienced setting up as much simpler than anticipated or more difficult than expected.

Of the 172 reviews, 71 of these explicitly refer to security, e.g., *“enter a name for the SSID and create a password, and that was it”* (R11280-R). The rest do not refer to security steps like setting up usernames and passwords as an explicit step in setup, e.g., *“Easy to setup and configure”* (R11107-R). Interestingly, almost all reviews that refer to security steps during setup are non-technical reviews (see Figure 2.2a).

More than half of the reviews that refer to security discuss problems with passwords, including one review which mentions having written the password on masking tape on top of the router. Of the rest, half a dozen were references to the relative ease of setting up using WPS and QRCode: *“Very easy setup using the WPS button on my router and the WPS button on the Extender”* (R12684-R). Others experience frustration with the same *“The robot will not scan the qr code when attempting to pair”* (R2976-S). Some of these reviews also mention returning devices because of friction during the setup process *“Every time, I tried to set it up, it said that it had failed!!!This router is getting returned!!!”* (R9692-R).

### FACTUAL INFORMATION AND ADVICE

In total, 62 (19.68%) reviews provided (purportedly) factual information, sometimes coupled with security advice to other users. Nearly half of these (43.5%) contain technical terms and details. Some of the technical reviews merely list the security protocols as part of the device specifications, e.g., *“The Range Extender supports n/g/b wireless with WEP, WPA/WPA2-PSK encryptions”* (R9436-R). This does not articulate whether these features are good or bad but may be useful for technically-literate consumers with matching expectations. Eight technical reviews draw some conclusions about the security of the device, but even these might be harder for a non-technical audience to grasp the implications, e.g., *“The continuous video on the SD card is accessible via the app (and the app servers are in the cloud like everything else) but is supposed to be end-to-end encrypted.”* (R467-S)

Nine reviews outline potential security issues in devices and contain technical advice on overcoming them. This is akin to ‘informal technical support’ normally provided by a ‘local expert’ to a device user [172, 180]. Four of these reviews ask users to isolate the devices on their network:

*“...Based on what I saw in the software I would want this camera completely isolated from the world. Don't use their app, don't scan the QR code, don't let it phone home to the internet. Put it on a completely isolated network with your NVR equipment...”* (R1654-S)

Of the reviews that contain security advice in non-technical language, two advise users to change the password of their devices *“I'm sure it's a default manufacturer password so I*

changed it, as I suggest anyone do" (R9144-R). Three warn users against buying products due to associated S&P issues:

*"Several stories in the media about the poor security of [device name] devices. If you value your privacy and security, and would prefer your personal data and camera feed information not to be sold to the highest bidder please avoid these cameras..." (R7518-S)*

A few of the non-technical reviews make a generic comment that a device is *secure* without providing any details, like in this cause for a travel router:

*"This is a great way to be more secure when using the Internet. Especially when traveling." (R12323-R)*

#### NEGATIVE OPINION ON SECURITY FEATURES

Unlike the previous theme, where users commented and advised on security features, the 49 reviews (15.6%) in this theme express strong negative sentiments about the perceived lack of security or privacy. This negative sentiment is about the security of the device as a whole and not about the steps involved in the configuration of security settings. The dissatisfaction includes general security concerns, notes about vulnerabilities, stories about devices being hacked, devices being flagged as non-secure, complaints about limited security, and discomfort with providing customer service remote access. That is, these are limitations of the security of the device, which in general cannot be remedied through any amount of configuration effort – the specification is unsatisfactory.

The security concerns raised vary based on the device type, and while some reviews raise these concerns in simple terms, others (21 of 49 reviews) use more technical language. For instance, both of the reviews below express concern about access to the video feed of a security camera, but the second is more technical in nature.

*"... Cons: [...] didnt even see an option to change username which is a big negative as these days privacy is of utmost importance..." (R5398-S)*

*"... your username:password credentials are passed in plaintext as part of the URL when interacting with the camera?..." (R5998-S)*

Security concerns expressed about routers include lack of encryption, packet sniffing, lack of an option to change the username 'admin', and vulnerabilities associated with the remote configuration of a router without a user account. We observed a difference again between technical and non-technical reviews, as with the following excerpts – the former describe the issue (R9257-R), and the latter merely states it (R12360-R):

*"The description says "Advanced Security" but it doesn't have WPA3 available nor is this device compatible with WPA3, maybe 10yrs ago it was "Advanced"..... There is vulnerability in WPA2..." (R9257-R)*

*"It has not increased the speed of signals.showing security problems" (R12360-R)*

Eight reviews provide accounts of devices the reviewer thought had been hacked; one was for a router, and the rest were surveillance systems. For instance:

*“... I logged in 1 morning (to the app) while hearing the clicks, and you can actually see a flash of light and a quick pause go off in the monitor with each click (as if someone is taking pictures).. it's as if this monitor's access is being live fed to perverted viewers who have access. Another thing I realised is that it only happened on the camera labeled 'bedroom'.” (R1738-S)*

Of note here is that some reviews situate S&P issues, where such ‘stories’ about security or privacy experiences have been seen to resonate with technology users in similar contexts [186], where stories are typically of perceived security incidents.

### USABILITY OF SECURITY FEATURES

Most (88%) of the 33 reviews within this theme are non-technical and can be broadly classified into security features hindering usability. A clarifying example of the former is a reviewer who is annoyed with 2FA authentication since it interferes with functionality:

*“Tonight the [device name] alarm went off whilst I was out. I accessed the app to see if I was about to be burgled or needed to speak to a visitor. And what did I get? ‘You need to setup 2 factor authentication before we will let you access your system’...” (R5901-S)*

Other reviews complain about the lack of support for password management software and talk about the difficulty in entering strong passwords:

*“...With the 4.0 line of firmware, stronger passwords are required. [...] It's also a bit annoying that you can't plug a standard keyboard into the USB port for typing the password. Or even use a touch screen (see above!) Using the mouse is not a friendly way to create a strong password, especially the tiny mouse the unit comes with...” (R2487-S)*

On the other end, we find reviews that comment on usability that is not security-enhancing, e.g., on how easy it was to connect to a router since there was no password required. One review notes a design complaint:

*“...‘I bought this router primarily for the WPA3 security. Problem is, when that's enabled, the onboard software disables the Wifi Protected Setup (WPS) button.’” (R9196-R)*

Another review complains about the lack of multi-user support:

*“... I tried to set up my wife with the app to access the camera and all she can do is view the camera/s or make them record, nothing else, which is not acceptable at all, every one I give access to must have the ability to do everything I can do...” (R7471-S)*

### POSITIVE OPINION ABOUT SECURITY FEATURES

The 28 reviews within this theme express a positive sentiment about the security features of the devices. Interestingly, the nine reviews in this theme that do not talk in technical terms offer details on why they like it—a handful mention encryption on the device, while others provide more details:

*“...I like the option of being able to share the video with other people [...] I wanted it to be secure in that no one was able to view it without my permission. [...] you can share the cam by sending a request direct to the other person's email, it also shows you on the app who has permission to view it [...]” (R4991-S)*

Several reviews indicate that users trust 2FA to be more secure and safe: *“The phone app also has 2 factor authentication (YEAH!! All apps should!! don't let the hackers into your IoT because they stole or guessed your password!)” (R467-S)*. The other reviews are satisfied with the security of the device because they trust their own configuration rather than the device itself:

*“...Personally I'm running super secure WiFi behind an awesome firewall and a VPN, there is noway some "hacker" is going to try that hard to get into this camera...” (R4409-S)*

The reviews for routers within this theme refer to the encryption settings, guest networks and built-in VPN, with one review mentioning that the built-in VPN was the reason for choosing a particular brand. In addition, a couple of reviews appreciated the DDoS and malware protection, firewall and networking monitoring tool that alerts them when a new device joins their network. Such reviews may indicate what options are available in the market for prospective buyers to then challenge where they see such options *not* being offered.

### DATA CAPITALISM AND PRIVACY

The 18 reviews in this theme are nearly evenly split between technical and non-technical reviews. Eight reviews express exhaustion at having to register with an account for usage and consent to user agreements:

*“...the app will not work at all unless you sign up for an account with [manufacturer name]. Not only is this completely unnecessary to operate the router [...] but I'm completely fed up with this behavior from companies. [...] I'm tired of "agreeing" that companies can do basically whatever they want without any legal repercussions or responsibilities...” (R10200-R)*

Another review talks about an app for an IP Camera (surveillance system) asking for permission to access the location and microphone of the users' phone, while another app requests permissions to tweet, comment and (un)follow on Twitter when logged in through a Twitter account. A couple of reviews embody resentment for having to sign up with an email identity to manage a router, and the underlying consensus across most of these reviews seems to be that their data is being 'sold' by the companies:

*“...The app asks for extra permissions like location data and even body biometrics. They are obviously selling this data...” (R1043-S)*

Research looking specifically at smart speakers [7] has found that users have insufficient understanding of the data that is collected and processed by these IoT devices. In this context, this translates to reviews demonstrating concerns about having insufficient information about privacy-related matters to have made an informed purchase decision. Prior research examining reviews [82] has surfaced consumer concerns about service practices, where we evidence specific features and stages in device use where consumers focus those concerns.

Conversely, when their privacy is protected, especially in surveillance systems, some users note and appreciate it. Two reviews refer to a privacy mode in their camera positively:

*“...What peaked my interest in this camera was the addition of face recognition, which means you are able to ignore family members if you want to. This was an important factor for me as the rest of the family were not too happy about being filmed all the time...” (R4721-S)*

However, one review mentions concern about how privacy is being handled:

*“...Security feels quite questionable. I don't see any promise your video/pics/data/camera is safe and secure. I can put a PIN on the camera and that's good, but other than that, I haven't noticed any real mention of how they are protecting my privacy...” (R4323-S)*

Aside from privacy issues, reviews generally indicate a weariness of having to create new, multiple accounts for a range of home IoT devices rather than this being consolidated. This then mirrors issues raised in password usability research.

## 2.5. DISCUSSION

The overarching result for RQ1 across all reviews analysed shows that, on average, 9.8% of reviews refer to S&P. There is some variance across device types, from surveillance systems with 5.9% to routers with 16.3%. Interestingly, once-vulnerable routers have a slightly lower percentage (13.6%), so the security problems that have plagued these devices have not emerged in Amazon reviews. Overall, it illustrates that a notable portion of reviews discuss issues related to S&P for these device types, compared to the 2% for home assistants [82]. This difference could be from our methodology - manual analysis allows for a more nuanced and reliable classification - or our device selection strategy; the market is relatively more mature for our devices than smart home assistants. Similar to S&P considerations historically being pushed to the end of the development life cycle or treated as an afterthought [206], S&P issues might permeate the market relatively late in the product diffusion curve [190].

Further, the results from the thematic analysis for RQ2 show that customers talk about their negative personal experiences with these devices, as evidenced by the reviews that discuss frustrations with device setup, stories of hacking, and exhaustion with what is perceived as unwarranted data collection from companies. Moreover, we also see reviews

providing security advice and voicing privacy concerns. Thus, our results indicate that at least a small percentage of consumers in the market for IoT devices *do* care about S&P and are able to articulate where these concerns arise in the lifecycle [126] of device ownership. There are then ‘signals’ in the market, the information within these reviews, that can be leveraged to ensure that other prospective buyers can make more informed decisions and to indicate to manufacturers that consumers care about S&P. This also signifies that amplifying this information would be useful on both these counts.

Looking further at RQ2, we find that some reviewers express concerns which inform a negative view of the device overall - their personal preferences were not met. This represents *dissatisfaction* with a newly-purchased device. Others express a range of positive and negative sentiments toward the setup phase for a newly-purchased device, where the most friction with security steps was experienced during setup (Fig 2.2b). This reveals the setup phase as a critical point in the ownership of a smart home device where information is needed.

The experiences of reviewers may point to a discounting of S&P risks in the market itself. It may be that a customer cannot return/exchange a device solely on the grounds of their S&P preferences not being met; we saw many reviews of owners being disappointed after purchase when new information about a device’s S&P properties becomes apparent to them, after some not-inconsequential period of attempted use; arguably those preferences are not being treated seriously at present, outside of regular no-questions-asked return policies. However, this points to a challenge for prospective buyers to *find* that information within reviews which may also discuss non-S&P issues and preferences. Even when they are able to find these reviews, the technical language used across most of the themes (Figure 2.2a) could hinder the usefulness of the information for those without a technical background.

A clear issue from a lot of the thematic analysis is that the customer immediately found that they did not get what they were expecting or, in time, uncovered an issue that they had not thought about, which then became a problem that they were stuck with. Any narrative of consumers choosing devices with inferior security must then also acknowledge those consumers who *know about* inferior S&P but do not know what to do with that new information (e.g., newly-discovered vulnerabilities, failings or oversights in S&P features).

### 2.5.1. LIMITATIONS

Our analysis focused on the specific subset of IoT device types that are commonly infected with Mirai. While this allowed us to determine consumer awareness of known S&P issues in the marketplace, they may not be representative of other IoT device types; many reviews did nonetheless highlight decision points where consumers had concerns (such as setup and determining S&P capabilities after purchase).

Our review sampling strategy does not optimise for any one ‘typical’ approach for review presentation or search; since Amazon presents product reviews in various ways, our sampling strategy was designed to best account for different review search strategies. In our analysis, we, for instance, uncovered specific S&P preferences, which can aid in accounting for more directed review search strategies.

Our analysis is limited to product reviews on an online marketplace. It is acknowl-

edged that the potential for various forms of media to inform security behaviours is under-researched (including TV and news [185, 187]); we add product reviews to that list, where here we found potential to leverage this additional source of information. Many of the reviews we analysed focused on the initial experience of device ownership, which presumably influenced the time when most users write reviews, that is, right after purchase. Where other research has recently begun to detail milestones in extended IoT device use (e.g., [46]), our analysis here pinpoints specific activities where S&P-issues were found, such as device setup.

### 2.5.2. RECOMMENDATIONS

Based on our analysis of customer reviews of home IoT devices, we provide the following initial recommendations:

- **Highlight S&P-related reviews.** Given that a not-inconsiderable portion of the reviewer population (9.8%, as in Table 2.3) articulates S&P-related concerns, it would be useful if e-commerce websites highlight these reviews. Irrespective of whether the reviews are fake or misinformed, highlighting the S&P signals in such reviews might encourage other buyers browsing through reviews to factor S&P prior to making a purchase decision. However, highlighting S&P reviews will only be partially successful since many of these (33.7%) are technical in nature and language. In order for it to be useful for a less knowledgeable buyer (Section 2.4.2), it would help to have an S&P specific rating for the devices. Such a rating will serve as a shorthand indicator of device quality and help consumers from a non-technical background interpret the sentiment expressed in these reviews.
- **Use the review system to match advice to emergent concerns.** Consumers may be willing to follow S&P advice if it addresses their existing concerns. The negative opinions examined in Section 2.4.2 surfaced S&P concerns and stories, for instance, and in Section 2.4.2, unusable security features were seen to hinder device use. If support can be provided to configure these features to the satisfaction of the user, it could reduce the lack of engagement with the devices; results in Section 2.4.2 highlight that some reviewers were able to find appreciable S&P features. On online shopping platforms such as Amazon, this could be addressed in the Q&A part of a product listing, with a variation of ‘signals of interest’ [191]. The market may not appreciate that a device has been bought but *later* unused due to S&P concerns – a further indication of ‘interest’ in using the device is then necessary to indicate that a solution was subsequently found. Where answers for S&P-related questions on the shopping platform can be indicated as useful, it can ‘match’ a solution to a concern and restore intentions to use a product beyond the initial signal of device purchase.
- **Design for shortcuts in security features.** We noticed a distinction between reviews framed in technical details, and others not (Section 2.4.2 and Figure 2.2a), with similar concerns around shared device use activities (such as setup, adding a new user, etc.). There is then scope to balance the needs of users who want detailed configuration options and novice users who want the ease of setup. This relates

to the established design principle of “Flexibility and efficiency of use” [166], and provides configuration options for both experienced and inexperienced users.

## 2.6. CONCLUSIONS

We investigated to what extent customer reviews of IoT products provide S&P information to consumers at the point of purchase. Where there were S&P signals in reviews, these included technical statements about features, frustrations with specific device use activities, as well as vignettes about trying to use a device in a particular context. Negative views on IoT devices were reflected in generally lower overall ratings for devices. All in all, we find that customer reviews provide a valuable and widely-used mechanism for conveying S&P information to consumers—prior to, and complementary with, potential future labelling schemes for IoT.

Our findings indicate that tangible options for S&P may be of interest as much as the features that participants can ‘imagine’, allowing users to compare meaningful options and offerings to choose from what is available rather than what is imaginable. This indicates that surveys of real device features in the market are useful. Future work will also include leveraging our manually-labelled reviews to train a classifier, to analyse a review dataset for S&P prevalence and themes.



# 3

## IIOT MARKET DYNAMICS

*We explore the relationship between the Security and Privacy (S&P) of IoT devices and their sales, considering the S&P signals in the context of these sales. We obtained expert S&P ratings of IoT devices from a European consumer association and the corresponding sales data from a leading Dutch online store. We complemented this with additional information like user ratings, the number of reviews and update support duration from two Dutch online stores. Our regression model shows that, holding other variables constant, a one-standard-deviation increase in S&P ratings corresponds to a noteworthy 56% boost in sales. Crucially, we observe a possible correlation between price and demand for S&P; at lower prices, the sales of IoT devices are directly proportional to the S&P rating, but this relationship diminishes as price increases. Further, we find that the presence of update support duration information, intended as a security signal, corresponds to higher S&P ratings and, all else being constant, also corresponds to a 69% increase in sales. While the exact causal mechanisms for the boost in sales remain unclear, our findings suggest positive incentives might be at play for IoT devices offering S&P at affordable prices and presenting relevant S&P information at the point of purchase.*

### 3.1. INTRODUCTION

A widespread consensus among security researchers states that many consumer Internet-of-Things (IoT) devices exhibit vulnerabilities that compromise user privacy and security. Despite the severe consequences of these vulnerabilities, ranging from individual breaches of privacy [53] to large-scale DDoS attacks [23, 124], there is no dent in the growth of IoT devices. The consumer IoT market is projected to grow steadily at 5% per annum from 2023 to 2028, reaching a market volume of around US\$232 billion by 2028 [205].

Observers argue that the sale of insecure IoT devices indicates market failures, namely externalities and information asymmetry [37, 56, 122, 197]. Externalities arise since manufacturers escape the consequences of poor security once the devices are in the market. Similarly, consumers might also undervalue security, since many of the negative

effects of their compromised devices end up with third parties, as in the case of DDoS attacks. Despite this lack of incentives, recent studies show that consumers not only care for IoT security and privacy, but they are also willing to pay more for it [70, 89]. However, it is unclear if these preferences for Security and Privacy (S&P) will translate into actual purchase decisions due to the lack of relevant S&P information during purchase. This information asymmetry implies consumers cannot factor S&P in their decisions, which in turn, implies devices with better S&P are not rewarded with higher sales. However, is this conjecture supported by empirical evidence?

In this study, we empirically examine this conjecture by conducting a comprehensive three-fold analysis of consumer purchase decisions on online stores. We analyse the purchase decisions using sales as a proxy, investigate information asymmetry in the context of these decisions, and explore whether presence of information intended to act as a security signal is associated with higher or lower sales.

Specifically, we first examine whether consumer preferences for IoT S&P are reflected in their real life purchase decisions. Using expert ratings of IoT S&P, including expert update ratings, obtained from the main consumer association in the Netherlands and the corresponding monthly IoT sales data from a leading Dutch online retailer, we answer the research question *'To what extent does the expert security and privacy rating of an IoT device correlate with its sales?'* Next, we study the information asymmetry in the context of these purchase decisions by evaluating whether information presented to consumers on online stores contain S&P signals. The information on online stores, like user product ratings, already serves as signals for device quality and vendor reputation [191] and influence purchase decisions [49, 61, 240]. Moreover, some user reviews on online stores also contain S&P related information [225]. In the absence of relevant S&P information, these signals of quality might act as proxies for S&P. To systematically examine if these signals are consistent with the S&P of an IoT device, we ask *'To what extent are expert ratings of S&P consistent with user ratings?'* We also analyze the presence of update support duration, presented as part of the product information on Dutch online stores due to a government mandate. It was intended to act as a security signal since IoT devices without updates can become unsafe [12]. We compare this to ratings from the consumer association asking *'To what extent are expert ratings of updates consistent with presence of update support information?'* Third, we also examine if the presence of update support duration is associated with a difference in sales through our final research question *'Does the availability of update support information of IoT devices on online stores correspond to higher sales?'*

To answer these research questions, we used three data sets. The Dutch consumer association, through its access to the technical labs commissioned by the European network of consumer associations, gave us access to expert ratings for four IoT device types: IP cameras, smart printers, smart speakers and smart watches. The expert ratings include the S&P rating, the update rating, and the overall device rating, an aggregate rating from various tests. From the leading Dutch online retailer mentioned earlier, we obtained the monthly sales data for the tested devices and the average prices they were sold at per month. The retailer opted to remain unnamed, and we will refer to them as "Winkel". Finally, we complemented the sales data with web scrapes of the product pages at Amazon.nl and at Winkel. We collected the average user rating, number of reviews

and product description including update support duration information for the tested devices.

To the best of our knowledge, this is the first study to use ground truth data on sales from a large e-commerce retailer. Although we focus on the Dutch market, prior research shows that IoT device popularity is similar across a variety of countries [132]. Moreover, we want to emphasise that this is an observational study. While we study the relationship between IoT device sales and other factors, we do not claim to establish causality. Rather, we want to test whether the observational data is consistent with our current notions of lack of reward for IoT devices with better S&P.

To answer our first research question, we construct regression models. We find that, after controlling for other factors like price of the device, a one standard deviation or 1.5 unit increase in the expert S&P rating corresponds to a remarkable 56% increase in sales. In contrast, an increase in the expert overall rating corresponds to only a 17% increase in sales. Moreover, we observe that the relationship between sales and expert S&P ratings is moderated by price. Across all IoT device types, at lower prices, higher S&P ratings correspond to higher sales, but this relationship diminishes as the price increases.

For our second research objective of evaluating information asymmetry, we find mixed results. We don't find any evidence of correlation between the expert S&P ratings and the user ratings. With respect to the update support duration information, we find that a mere 4.3% of the product links under consideration contain complete update support duration information on both the online stores, while 27.8% of the links contain complete information on at least one of the stores. However, our findings indicate that the few IoT devices with complete update support information on both Amazon and Winkel exhibit a significantly higher expert rating for updates. This suggests that availability of update support duration information on online stores is reflective of the general update practices of manufacturers and can reduce information asymmetry.

For our third research question on the difference in sales due to presence of update support information, we extend the regression model, built to answer our first research question, with update information status on Amazon.nl and Winkel. We find that, after controlling for the other factors, devices with complete update support information on Winkel correspond to 69% higher sales. This suggests that there are some positive incentives at play for sellers and manufacturers that provide the additional information although the exact mechanism of the incentive is beyond the scope of our research. Our empirical study thus provides a nuanced perspective on IoT device interactions driven by real-world sales data and expert IoT ratings. Theoretically, our findings help understand how barriers to security from economic theory, like information asymmetry and externalities, play out in real markets.

Finally, although tangential to our research goals, our data allows us to capture one additional scientific benefit: we can quantify to what extent public data from online stores can be a proxy for ground truth on sales, since we obtained the latter from Winkel. We find that the number of reviews on Winkel is a reasonable proxy for the sales data from Winkel. In contrast, the number of reviews on Amazon.nl showed low correlations with the Winkel sales data. This suggests that although the proxy is useful, it specific to the retailer.

The paper is structured as follows: In Section 3.2, we review related research on

consumers' IoT purchase decisions and information asymmetry. Section 3.3 provides details on the datasets used. Section 3.4 explores the relationship between expert S&P ratings and sales, answering our first research question, while Section 3.5 and Section 3.6 address our second and third research question respectively. Section 3.7 discusses our findings and offers recommendations, and Section 3.8 concludes the paper.

## 3.2. RELATED WORK

In this section, we elaborate on related studies on various aspects of consumer purchase decisions in the market for IoT devices. We also present other work that evaluated the security and privacy of IoT devices using different techniques.

### 3.2.1. CONSUMER PURCHASE DECISIONS AND SIGNALS FOR S&P

There has been a fair amount of prior work that studied consumer purchase decisions on online stores. Several studies have found evidence that high user ratings on online stores correlate with higher sales of products [49, 88, 153, 240]. Similarly, several studies have noted that the volume of reviews on the platforms also influence sales [17, 61, 141, 199].

With regard to S&P and purchase decisions, interviews with retail customers highlighted that the point of purchase is an opportune moment for informing consumers about the security of the devices purchased [177]. Other studies delved deeper into this in the IoT context, analysing purchase decisions when adequate information about the S&P of the device is provided. Blythe, Johnson, and Manning [34] find evidence that providing clear security information before purchase can promote secure device selection. Gopavaram et al. [90] found that consumers were willing to pay more for privacy when there were clear indicators of the privacy level. Emami-Naeini et al. [70] empirically confirmed these findings in an incentive compatible user study and showed that consumers are willing to pay more for S&P when relevant information is presented.

There has also been prior work that evaluated different modes of presentation of S&P. Various forms of information labels have been proposed including graded labels, labels denoting S&P features, and labels indicating approval through independent assessment [111]. Emami-Naeini et al. [72] received positive feedback for labels with a similar design to food nutrition labels. Morgner et al. [157] found support for provisioning update support duration information, particularly from those perceiving higher risks from use of IoT devices. Although security labels have not yet found widespread adoption in the market, display of update support duration information has been in effect on Dutch online stores since 2020.

Based on these studies, we consider average user ratings, number of reviews and update support duration information as possible variables that influence purchase decisions. Moreover, to increase the generalizability of our findings, we collect these variables from two popular online stores in the Netherlands – Amazon.nl and Winkel.

### 3.2.2. ESTIMATES OF IoT DEVICE POPULARITY

Earlier studies have estimated the popular devices in the market, which can be taken as an indicator of their sales, in a variety of ways. These studies are typically positioned as device identification in the wild and they use various techniques to identify the brand, make

and model of deployed IoT devices. Several of these use network traffic to identify IoT devices [152, 194, 200, 234], however, we do not include them here since these analysis are local to the network from which the measurement was done and cannot be generalized. Other techniques that allow for more generalizability include passive sampling of network scans and active experiments done as ISPs and IXPs [192] and DNS fingerprints [178]. Taking a different approach, Kumar et al. [132] conducted a large-scale empirical study of IoT devices installed in homes of users of Avast antivirus software. They show that barring a few exceptions, the popularity of the four IoT device types that we consider – IP cameras, smart printers, smart speakers and smart watches – are largely similar across geographies.

### 3.2.3. ESTIMATES OF SECURITY OF POPULAR IOT DEVICES

Junior et al. [114] estimated the security of popular IoT devices by analyzing the vulnerabilities present in their companion apps. For device selection, they started with the 100 most popular smart hubless devices on Amazon and then filtered it to devices that use WiFi for communication and then further to the categories of smart plugs, bulbs and infrared controllers. They find many devices share a common app and that 50% of apps corresponding to 38% of devices do not use proper encryption techniques. Moreover, there are papers [203, 238] that evaluate the security of popular IoT devices without specifying the metric used to define popularity for the device selection which lends their results incomparable with ours. In our work, we use sales figures for popularity and estimates of device security and privacy from the tests conducted by commissioned labs across Europe.

### 3.2.4. CORRELATION BETWEEN S&P AND SALES

To the best of our knowledge, there are no studies that estimate the relationship between IoT device security and privacy, and its sales. However, earlier studies have analyzed the relationship at the level of a firm – between a firm’s security and privacy posture and its stock market evaluation. Boroomand et al. [35] found that as a firm’s investment in data privacy increases, its market valuation decreases. They attribute this to information asymmetry – firms’ are not able to communicate their increased investment in privacy and consumers are not able to reward it despite valuing privacy. A similar study [237] evaluated the effect of investment in Data Privacy and Security (DPS) on the market value of two kinds of firms – those dealing with Business Data Analytics (BDA) and those that do not. They found DPS investment decreases a firms systematic risk and that these effects are higher for non-BDA firms when compared to BDA firms.

## 3.3. DATASET DESCRIPTION

In this section, we describe the three datasets from four sources that we used in our study. First, we received expert IoT device ratings from *Consumentenbond* (CB), a consumer welfare organization in the Netherlands. Next, the ground truth of sales data of IoT devices was generously provided by ‘Winkel,’ a leading online retailer in the Netherlands. Both of these datasets contain competitive data and are therefore typically confidential, sensitive, and strategically significant, which often restricts their availability for research. However,

Consumentenbond and Winkel were exceptionally generous in providing these datasets for our study. Finally, we collected publicly visible attributes on online stores that prior work (Section 3.2) has identified as influencers in purchase decisions: user rating, number of reviews, and information on update support duration. We collected these from the IoT device product pages of two sources, the websites of Winkel and Amazon.nl, the Amazon store for the Netherlands. Table 3.1 provides an overview of the sources and the data obtained from each of them. In the following sections we describe these data sets. We use the term user ratings to refer to the average user ratings on Amazon and Winkel, and expert ratings to refer to the CB ratings.

**Table 3.1:** Overview of the collected data.

Dataset	Data Source	Data obtained
Expert IoT device ratings	Consumentenbond (CB)	Security and privacy ratings Update ratings Overall device rating
IoT sales data	Winkel	Monthly sales figures Average price per month
Information from product pages	Amazon and Winkel	User rating Number of reviews Update support duration

### 3.3.1. EXPERT IOT DEVICE RATINGS FROM CONSUMENTENBOND

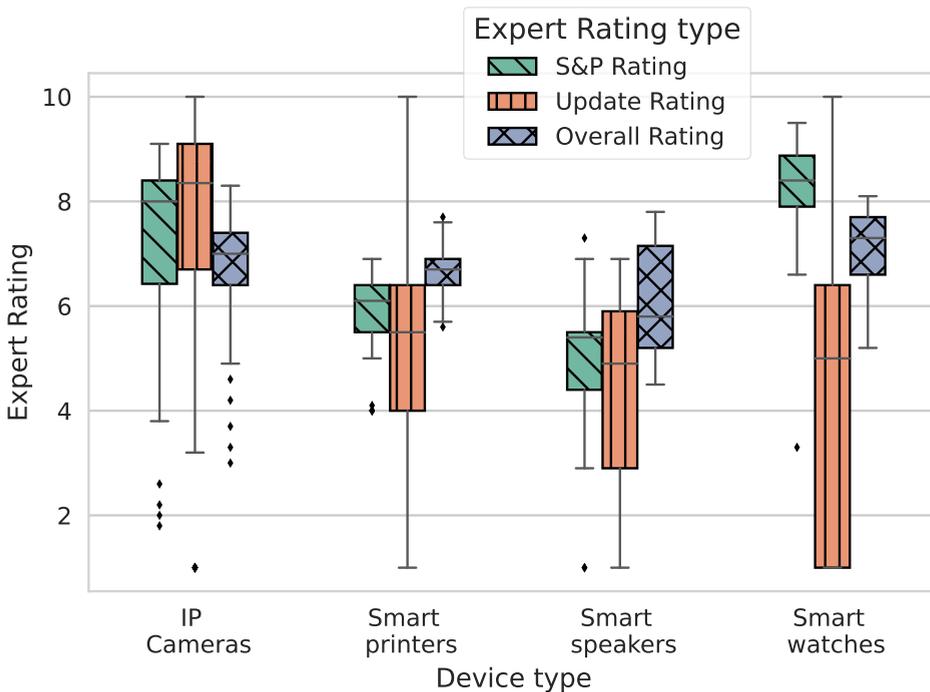
The Consumentenbond is a non-profit consumer welfare organization in the Netherlands.<sup>1</sup> It works on a membership model and currently has about 420,000 members (about 5% of all Dutch households). Through a network of professional technical test labs commissioned by various European consumer associations, they conduct independent tests of various devices and provide valuable advice to consumers to enable informed purchases.

We received all the test results for the four IoT devices types tested by CB - IP Cameras, smart printers, smart speakers, and smart watches. All the tests were conducted between 2016, the start of IoT device testing by CB, and 2023. The testing process remained largely stable over this time period, with most devices being tested soon after their release date. In some cases, often due to the devices gaining popularity after their initial release, there were delays between device release date and the testing date.

The tests conducted and ratings given by CB are exhaustive and span a broad range, however, we only used the three expert ratings that were relevant for our study – the S&P rating, the update specific rating and the overall device rating. The S&P rating is an aggregate of the results of different sub-tests evaluating the password policy, number of vulnerabilities, updates and privacy concerns like data sharing. The heterogeneity in the device types, in terms of attack surface, privacy concerns and device functionality, leads to the necessity of device specific tests. Thus, the S&P sub-tests per device type vary both in number and type. Unfortunately, our confidentiality agreement with CB prevents us from sharing the test specifics and details on how the S&P rating is derived, but for each

<sup>1</sup><https://www.consumentenbond.nl>

device type, the areas covered by the sub-tests are added in Appendix B.1. The update rating captures the clarity about update availability, and in the case of smart speakers, also the automatic update options available on the companion mobile apps. Although the update rating is included in the S&P rating, we also analyze it independently, as updates represent one of the most visible aspects of manufacturers' commitment to security from a consumer perspective. The overall rating is an aggregate of the S&P rating, update rating and other ratings like ease of use, and some device specific ratings like sound quality. Figure 3.1 shows the distribution of all three expert ratings across the different device types, the ratings are between 1 and 10. Although the S&P ratings are skewed towards higher values in some cases, we find enough variation among the ratings for a meaningful analysis of the relationship between S&P and market performance.



**Figure 3.1:** Distribution of expert ratings across the four device types

In total, we had the expert ratings for 469 IoT devices (130 IP cameras, 174 smart printers, 31 smart speakers and 134 smart watches) tested by the Consumentenbond. The average S&P rating across the four device types is 6.9 with the lowest being 1 for a smart speaker and the highest 9.5 for a smart watch. As seen in Figure 3.1, even within the same device type, there is a wide spread in the S&P ratings. This suggests that despite the potential disincentives for manufacturers to focus on S&P, at least some IoT devices have better S&P than others. We also observe variance across the four device types. The average S&P rating is the lowest for smart speakers (4.8) followed by smart printers (6) and

IP cameras (7.3), with smart watches having the highest average S&P rating (8.3). The low average S&P rating for smart speakers is primarily due to the low scores of the companion mobile apps. It is important to note that these ratings are from the device population that was tested by CB and cannot be generalized to the device types.

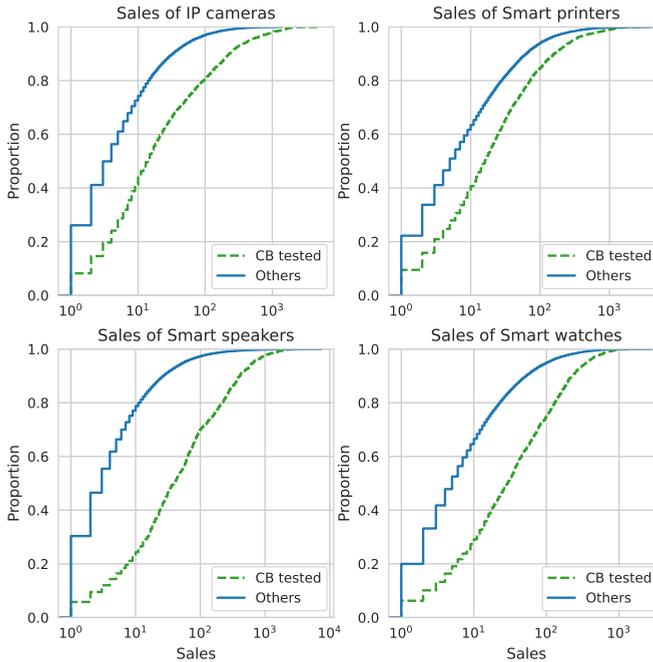
The average update rating across all devices (5.2) is lower than the S&P rating (6.9). Smart watches have the lowest average rating for updates (4) closely followed by smart speakers (4.4) and smart printers (5), while IP cameras have the highest average of 7 for updates. The lower average update rating across IoT devices aligns with findings from other studies that indicate that update deployment tends to be slower for IoT devices [181]. The overall rating has an average of 6.8 across the device types with the lowest being 3 for an IP camera, and the highest 8.1 for a smart watch. In contrast to S&P rating and update rating, the average overall ratings across the device types do not show much variation.

### 3.3.2. SALES DATA OF IOT DEVICES

As mentioned earlier, we got sales data of IoT devices from a leading online retailer in the Netherlands. This was for a period of almost four years, from January 2019 to August 2023, and contained two datasets. The first dataset had the sales figures and average prices of IoT devices that were tested by CB. We used European Article Number (EAN), a unique 13 digit identifier for products, to match the CB test results with the Winkel sales data. The unit for the sales figures is the number of items of each device that were ordered and also reached the customer. Like most online retailers, Winkel also uses dynamic prices that are determined by a range of factors. For the purpose of this study, along with the monthly sales data, we obtained the monthly average prices at which the devices were sold for 385 devices (113 IP cameras, 132 smart printers, 28 smart speakers and 112 smart watches) that were tested by CB. We were not able to obtain the sales data for 18% of the devices tested by CB either due to the devices not being sold on Winkel or because we did not have the EAN code to match the devices with the sales data.

The second dataset had the sales figures and average prices of IoT devices within the same device type category that were not part of CB's testing. Comparing these two allowed us assess the selection bias in the devices tested by CB. To ensure the tests are relevant for a larger customer base, CB uses third party market research data to select and evaluate the more popular devices. When we compare tested vs. non-tested devices, we can indeed see that the tested devices have higher sales. [Figure 3.2](#) shows the cumulative distribution of sales over the years of devices tested by CB and those not.

Across all the device types, the average sales per month of devices tested by CB is 91.36 (min: 1, median: 19, max: 5,353) while its 20.7 (min: 1, median: 4, max: 7,394) for others that were not in the CB data set. The comparison also underscores the rationale behind CB's choice to prioritize testing popular devices. As observed, approximately 20% or more of the devices across all the device types that CB has not tested have recorded only a single sale. We also find that the devices tested by CB are marginally more expensive than the other devices, especially for smart speakers and watches. The ECDF of the price comparison is added in the [Appendix B.2](#).



**Figure 3.2:** CDF of sales of devices tested by the Consumentenbond and those not

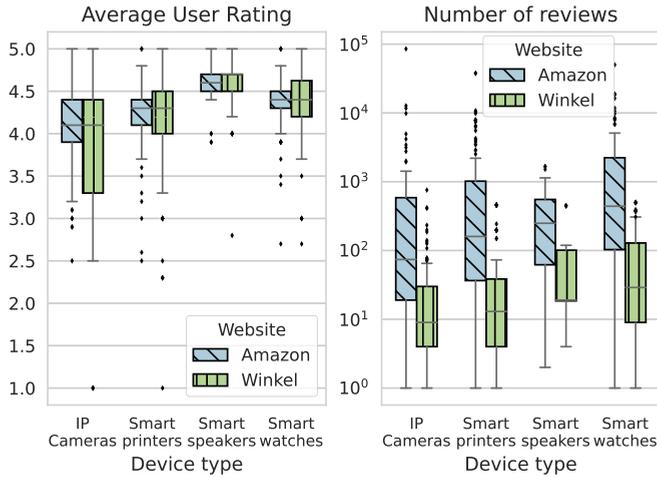
### CORROBORATING WINKEL SALES DATA WITH COUNTRY LEVEL SALES

To ensure that the sales data obtained from Winkel is representative of the national trends, we compared it to a commercially available Dutch market study data that is occasionally used by CB for device selection. Although we did not have this data for smart printers and speakers, for IP cameras and smart watches, a Spearman's Rank correlation test revealed a statistically significant and high positive correlation between the two data sets (0.72 and 0.75). This suggests that although there might be slight differences in the sales across specific stores, the sales data from Winkel is reflective of the national trends in sales of IoT devices.

### 3.3.3. INFORMATION FROM PRODUCT PAGES

As noted earlier, we complemented the sales data with information scraped from two popular online stores, including user ratings, the number of reviews, and the duration of update support. We collected this data by scraping the product pages of the devices under consideration on Amazon.nl and Winkel. The preliminary collection was done between May and June 2023, with some missing details added in July 2023. Out of the 469 devices with expert S&P ratings, there were 14 devices that were not found on either platform. Of the remaining 455 devices, at the time of collection, 335 were found on both, 32 were found only on Amazon<sup>2</sup> and 88 were found only on Winkel.

<sup>2</sup>Unless otherwise specified, all references to Amazon are limited to Amazon.nl



**Figure 3.3:** Distribution of user ratings and total reviews on Amazon

### USER RATINGS AND NUMBER OF REVIEWS

In total, we scraped the details of 367 devices from Amazon, across 455 product links (122 IP cameras, 170 smart printers, 39 smart speakers, and 124 smart watches). On Winkel, we scraped details of 423 devices across 505 product links (140 IP cameras, 185 smart printers, 45 smart speakers, and 135 smart watches). On both Amazon and Winkel, the number of product links is higher than the number of devices because, for some devices, we found multiple product links due to variations in colours, sellers, or device bundles. In all such cases, we collected information from all the product links and treated them as separate instances rather than aggregating them per device to allow for a detailed analysis of update support information.

Figure 3.3 shows the distribution of user ratings and number of reviews on both the platforms. We find that the mean of user ratings are comparable across both the websites, 4.26 on Amazon and 4.17 on Winkel. However, the differences in number of reviews is staggering. The average number of reviews on Amazon is 1792, while it is 64.85 on Winkel, a mere 3.6% of the average number of reviews on Amazon. This is likely due to Amazon's policy of aggregating reviews across all of its websites rather than a difference in popularity between Amazon.nl and Winkel. For instance, Amazon.nl also has reviews from Amazon.com, the US Amazon store and Amazon.de, the German Amazon store.

### UPDATE SUPPORT INFORMATION

As mentioned in the Introduction, after a government intervention, sellers agreed to provide consumers with information about the availability of updates for smart devices at the time of purchase [12]. Note that there is a distinction here between seller and manufacturer. The seller acts as an intermediary between the manufacturer and consumer, listing the products on the e-commerce store and managing the sales. As per the intervention, manufacturers must provide the information to the sellers, who, in turn, update it on the online store to inform consumers. This was agreed upon in 2020. We started our

data collection in May 2023, approximately three years after this policy came into effect allowing us to also analyse the presence of update support duration information as a security signal.

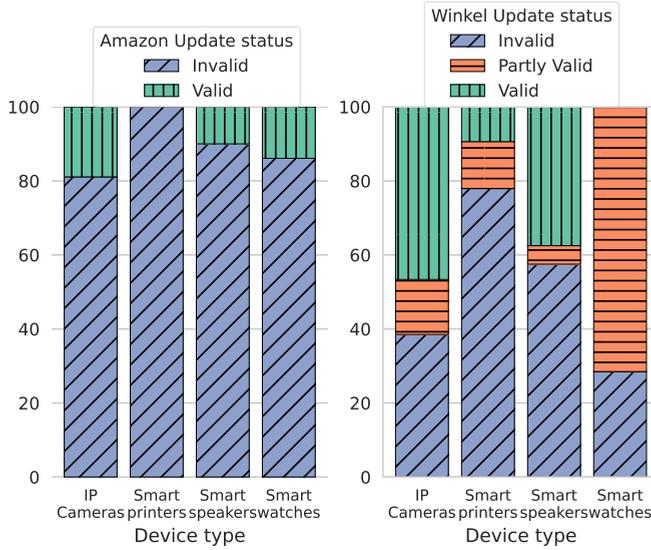
Our study reveals differences in how the update support duration information is displayed on each website. On Amazon, the update information is presented as a date (e.g., '13 April 2030'), and the update support field is part of the product's technical specifications. On Winkel, the update availability is specified using three fields under 'Introduction and support' within the product specifications table. These three fields are *Introduction year*, *Introduction month*, and *Support with updates*. The update support is specified in terms of the number of months from introduction (e.g., 'At least 24 months after the date of introduction'). While this clarifies the manufacturer's commitment to update support, some products have incomplete information as not all three fields are filled out.

Of the 455 products on Amazon, only 54 have a valid date; for the majority, the field is either empty or says 'Unavailable.' We categorise both empty and unavailable update information as invalid. Moreover, for six of the 54 products, the date is in the past (e.g., '30 June 2022'), and for one, it says ('1 January 2099'). However, we conservatively count these as valid since it is beyond this research's scope to verify if the updates for these devices continued past 2022 or will continue till 2099. Of the 505 products on Winkel, 231 had no information on updates whatsoever, 113 had valid information with all three fields complete, and 161 had partly valid information. That is, there was information about the duration of update support, but the Introduction year or month fields were empty or vice versa. We tag these as 'partly valid.'

Figure 3.4 shows the distribution of update statuses across the device types. On Amazon, we find that none of the smart printers have any valid information about updates; among the other device types, there is little variation – they all contain similarly minuscule percentages of valid information. Similarly, on Winkel, smart printers have the highest percentage of invalid information (78%). However, around 46.6% of IP cameras and 37.5% of smart speakers have valid information, while 71.5% of smart watches contain partly valid information – they provide information on the updates but fail to mention additional information on introduction year and month.

#### COMPARISON OF UPDATE SUPPORT INFORMATION ON AMAZON AND WINKEL

We find that for the 335 devices across 417 product links available on both websites, there is substantial divergence in the update statuses. Table 3.2 provides a comparison of the update field information. Only 18 devices, a meagre 4.3% of devices that are available on both platforms, provide valid update information on both websites, while an additional 27.8% provide valid information on one website but not on the other (7.9% have valid information on Amazon but not on Winkel and 19.9% on Winkel but not on Amazon). This suggests that the provisioning of update information is dependent on seller practices rather than the device-specific update support information provided by the manufacturer. Further, for none of the 18 devices with valid update statuses on both websites does the update support duration match. On Amazon, 14 of the 18 devices mention update support until 'April 2030,' whereas on Winkel, most devices mention update support for 24 months after introduction, with varying introduction years and months. This further



**Figure 3.4:** Distribution of Amazon and Winkel update statuses across device types

highlights the role of sellers in provisioning this information and the difficulty in ensuring accuracy of such information.

**Table 3.2:** Comparison of update status between Amazon and Winkel for devices found on both

Amazon update status	Winkel update status	Count
Invalid	Invalid	163 (39.1%)
Invalid	Partly Valid	120 (28.8%)
Invalid	Valid	83 (19.9%)
Valid	Invalid	11 (2.6%)
Valid	Partly Valid	22 (5.3%)
Valid	Valid	18 (4.3%)

### 3.4. RELATIONSHIP BETWEEN S&P AND SALES

In this section, we analyse the datasets described so far to answer our first research question on the relationship between security and privacy of an IoT device and consumer purchase decisions with sales as a proxy. To achieve a holistic understanding of the interplay among various factors, we also consider the variables previously identified in the literature as influencers on sales, rather than restricting ourselves solely to S&P. Moreover, in order to comprehensively assess the interplay and associations of the different variables with the sales, we used a Generalised Linear Regression (GLM) model modelled using the `glmmTMB` package in R [41]. Since our dependent variable, the sales of IoT devices, is a count data, we had two options for the GLM – Negative Binomial and Poisson Models.

The sales data exhibited over dispersion, meaning the variance was larger than the mean. Therefore, we used the Negative Binomial model since it accommodates excess variability. Moreover, to generalise interactions among variables across different device types, we employed a Mixed Effects model. This approach treats device type as a random effect and considers other variables as fixed effects, allowing us to derive broad-level inferences independent of device-specific variations in the test protocols.

### 3.4.1. EXPLANATORY VARIABLE SELECTION

Our explanatory variables of interest are price, the expert S&P, update and overall ratings from CB, user rating and number of reviews on Amazon and Winkel. Prior to running the model, we conducted multicollinearity tests to remove explanatory variables that are highly correlated with each other. This ensures stability and reliability of our parameter estimates by mitigating issues that arise from high correlations among explanatory variables. We used Generalised Variance Inflation Factor (GVIF) to check for multicollinearity. Following suggested guidelines, we excluded the expert update rating from our analysis due to its GVIF value exceeding 5 [173]. This left us with eight remaining variables, Device type, price, expert S&P rating, and Overall rating from CB, user rating and number of reviews from Amazon and Winkel. The GVIF of the variables included in the model is added in Appendix B.3.

Moreover, we conducted bi-variate correlation tests between the dependent variable and the explanatory variables to get a sense of how each explanatory variable relates to the dependent variable, and validate its inclusion in the model. To control for multiple comparisons, we adjusted our p-values using the False Detection Rate (FDR) approach, implemented through the Benjamini & Hochberg method in the stats package in R [184]. There were no changes in the significance levels of the correlation tests after the adjustment. We explore these correlations briefly before analysing the results of the model.

#### CORRELATION BETWEEN SALES AND PRICE

We first analyse the correlation between the price and sales to better understand the price sensitivity of IoT devices. Since neither data followed a normal distribution, we used Spearman's rank correlation test. The results, as shown in Table 3.3, reveal a statistically significant weak negative correlation between price and sales for IP cameras and smart printers and a moderate positive correlation for smart speakers. This suggests that consumers possibly perceive IP cameras and smart printers as utilitarian, where price and affordability plays a decisive role in their purchase decisions. This observation may also be influenced by the competitive dynamics in these markets, where companies often compete on price. In contrast, for smart speakers which often come with advanced features and innovations, consumers might be more willing to pay a premium for additional features. Moreover, smart speakers are often part of a larger ecosystem and consumers invested in these ecosystems might be willing to pay more for compatibility.

#### CORRELATION BETWEEN SALES AND EXPERT RATINGS FROM CB

Next, to understand if and to what extent S&P features and the overall device quality as indicated by the CB expert ratings correspond with the sales of a device, we performed pair-wise correlations between each of them (Table 3.4). We found no significant correlations between the expert S&P rating of a device and the sales. This implies that a direct

**Table 3.3:** Correlation between sales and price from Winkel

Device type	Spearman correlation
	Total sales vs. average price
IP Cameras	-0.251**
Smart printers	-0.339***
Smart speakers	0.517***
Smart watches	0.071
Aggregate	0.006

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

linear relationship between a device's S&P posture and its sales is not readily evident. We find statistically significant correlations between sales and expert overall rating across all devices and also at the aggregate level. Due to the significant correlations between the expert overall ratings and sales, we include it in the model. Moreover, despite the lack of correlation, we also include the expert S&P rating to gain a deeper understanding of the underlying relationship between S&P and sales which might be more complex than a linear association.

**Table 3.4:** Correlation between sales and CB ratings

Device type	Spearman correlation	
	Total sales vs	
	Expert S&P rating	Expert Overall rating
IP Cameras	0.088	0.216**
Smart printers	-0.035	-0.286***
Smart speakers	0.100	0.358**
Smart watches	0.037	0.333***
Aggregate	0.052	0.195***

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

### CORRELATION BETWEEN PRICE AND EXPERT RATINGS FROM CB

Next, we explore the relationship between the average price and the expert CB ratings. As seen in Table 3.5, the price of IP cameras and smart watches exhibit a modest positive correlation with the expert S&P rating. In contrast, for smart speakers, we find a moderate negative correlation between the price and S&P rating. However, smart speakers are a unique case. Within the 27 smart speakers in our sample, four smart speakers with the highest sales collectively account for approximately 56% of all smart speaker sales. These four devices have prices higher than average and expert S&P ratings lower than average. This indicates that the negative correlation is possibly influenced by the disproportionate presence of these devices, and cannot be generalised.

With respect to expert overall rating, we find moderate positive correlations between the sales and average price at the aggregate device level and also at the individual device level for all devices except IP cameras. The lack of correlation for IP cameras might be due to various factors like a more mature market, diverse manufacturers, differentiated

market segments based on price and features, or other market factors. Nonetheless, the high degree of associations between the average price and expert CB ratings for the other device types suggests that these two variables exhibit interaction, which could, in turn, influence sales. To accommodate this interaction, we incorporate an interaction term between price and expert CB ratings in the model.

**Table 3.5:** Correlation between average device price and expert CB ratings

Device type	Spearman correlation	
	Average price vs	
	Expert S&P rating	Expert Overall rating
IP Cameras	0.259**	0.160
Smart printers	0.112	0.633***
Smart speakers	-0.486***	0.756***
Smart watches	0.429***	0.559***
Aggregate	0.092	0.44***

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

#### CORRELATION BETWEEN SALES AND USER RATING AND NUMBER OF REVIEWS ON AMAZON AND WINKEL

We find many significant positive correlations between sales and both user ratings and number of reviews from Amazon and Winkel (Table 3.6), validating their inclusion in the model.

Moreover, this analysis allows us to determine to what extent publicly available information like number of reviews can serve as a proxy for sales. We find that the sales of IoT devices show a statistically significant positive correlation with the number of reviews on both Amazon and Winkel. This is true at both the aggregate level, across all the device types, and also at the individual device level. The only exception is the correlation between the sales of IP cameras and the number of reviews on Amazon.

Furthermore, we find that the correlation between the sales and the number of reviews on Winkel (0.8) is much higher than on Amazon (0.259). This implies that the number of reviews on a platform can act as a reliable indicator for sales on that platform. Since Amazon aggregates its review count across all of its stores, the number of reviews on Amazon may not be a reliable proxy for sales.

In contrast, with respect to the user rating, the reverse is true. The correlation between the sales and user rating on Amazon (0.415) is higher than the correlation with the user rating on Winkel (0.193). This suggests that even on the same platform, user rating is a poor indicator of popularity. However, since, similar to the number of reviews, the user ratings on Amazon are aggregated from different stores like Amazon.de, Amazon.com, etc., the Amazon ratings seem to be a marginally more reliable indicator of popularity. It is worth noting that the devices in our data set have a selection bias of being popular in the Netherlands and Europe, and within this set of devices Amazon user rating shows a moderate correlation with sales. Without further research, we cannot conclude if the finding will hold true in other geographies or for other devices.

**Table 3.6:** Correlation between total device sales and consumer metrics from Amazon and Winkel

Device type	Spearman correlation			
	Total sales vs			
	Amazon user rating	Winkel user rating	Amazon total reviews	Winkel total reviews
IP Cameras	0.529***	0.439***	0.094	0.762***
Smart printers	0.189	-0.001	0.278***	0.777***
Smart speakers	0.488**	0.397**	0.376**	0.651***
Smart watches	0.213**	-0.191*	0.338***	0.767***
Aggregate	0.415***	0.193***	0.259***	0.802***

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

### 3.4.2. MODEL EVALUATION

Next, we outline the steps involved in Modelling. As mentioned earlier, we use a Mixed Effects Negative Binomial Model, treating the explanatory variables as fixed effects and the device type as a random effect. This helps us understand the overall trends that affect all devices while also capturing device type specific variations. To validate our decision to include the device type as a random effect, we compared the fit of two models using likelihood ratio test [137]. The first model contained device type as a fixed effect, and the second as a random effect. The p-value from the likelihood ratio test was significant ( $p < 0.001$ ), indicating that the model with device type as a random effect was a better fit.

Moreover, as a preprocessing step, we scaled and centered the variables to ensure uniform influence and improve model convergence and interpretability. Scaling adjusts variables to have similar scales, preventing larger values from dominating. Centering sets variable means to zero, easing interpretation and providing a meaningful intercept.

We added the eight variables identified in the previous section in a step-wise forward manner. We explored different ways of ordering the fixed effect variables during the step-wise forward training but found that the ordering does not make any difference to the results. We therefore picked an intuitive ordering for ease of presentation. In the model presented, we start with the intercept only model, then add the device type, the average price, the expert S&P rating and the expert overall device rating, followed by user ratings and number of reviews on Winkel and Amazon. Further, as mentioned earlier, we added interaction terms between price, expert S&P rating and expert overall rating of CB. This enabled us to evaluate how the relationship between these variables and sales varies with changes in the other variables. This resulted in eight distinct models, which are shown in ?? in Appendix ??.

To evaluate which model is best suited for our data, we checked the goodness of fit using Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC) and Log-likelihood, following best practices from literature [42, 65, 100, 120]. AIC balances fit and complexity, guarding against over fitting, while BIC penalises complexity, promoting simpler models. Lower AIC and BIC values signify better model fit [42, 65, 120]. Log-likelihood quantifies how well a statistical model explains observed data, with higher values indicating better fit [100]. Moreover, in a mixed effects model, two  $R^2$  types are used to assess the model: Conditional  $R^2$ , which gauges the explanatory power of the

fixed effects and Marginal  $R^2$ , which measures the combined explanatory power of both fixed and random effects. Higher values for both indicate improved model fit.

Based on these criterion, amongst the models in ?? (Appendix ??), we identified Model 8 to be the best fit for our data. It has a lower AIC, BIC and higher log likelihood. Moreover, although the conditional  $R^2$  is slightly lower than the previous model, it has the highest marginal  $R^2$  amongst all the models. This indicates that the entire model, including the fixed and random effects, best explains the variation in our dependent variable, the total sales.

### 3.4.3. MODEL INTERPRETATION

The Incidence Rate Ratios (IRR) of each of the explanatory variables in our model is shown in Figure 3.5. The IRR quantifies the multiplicative change in the rate of occurrence of an outcome (in our case, the sales) when the explanatory variable changes by one standard deviation, after controlling for the other variables. An IRR of 1 suggests no change, while greater and lesser than one denotes an increase and decrease in sales respectively. The interaction terms added between price, expert S&P rating and expert overall rating allows us to see the combined effect of these variables on the sales. Table 3.7 shows the mean and standard deviation of all the explanatory variables in the model.

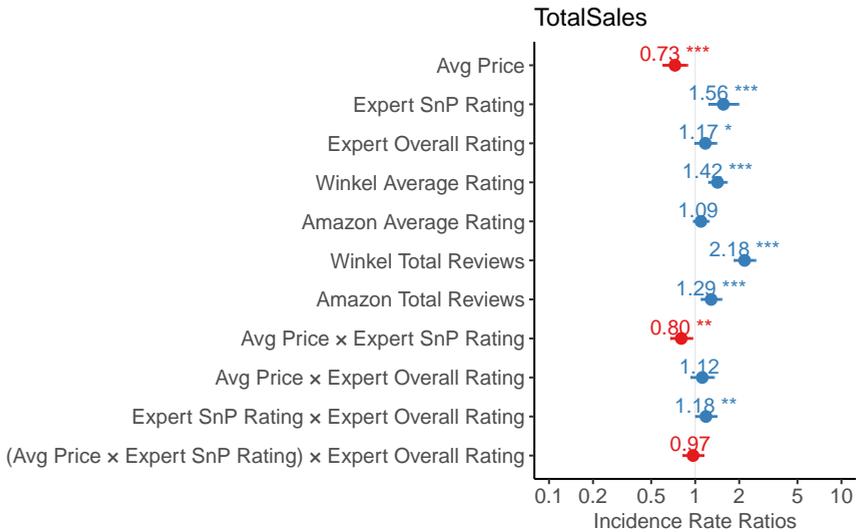
**Table 3.7:** Mean and Standard Deviation of Explanatory Variables

Explanatory Variable	Mean	Standard Deviation
Avg Price	€179.10	€112.70
Expert S&P Rating	7.05	1.48
Expert Overall Rating	6.94	0.65
Winkel User Rating	4.18	0.64
Amazon User Rating	4.27	0.38
Winkel Total Reviews	67.21	120.11
Amazon Total Reviews	1966.33	6665.94

To address our first research question regarding the relationship between expert S&P ratings and consumer purchase decisions of IoT devices, we found that, when holding other variables constant at their mean values, an increase in the expert S&P rating by one standard deviation accounts for a 56% increase in sales (IRR of 1.56). In contrast, an increase in the expert overall rating by one standard deviation corresponds to only a 17% increase in sales (IRR of 1.17). Additionally, a one standard deviation increase in price corresponds to a 27% decrease in sales (IRR 0.73). Moreover, the significant interaction term between price and expert S&P rating indicates that the interaction between them dampens the effect of either variable on sales by 20% (IRR 0.8). On the other hand, the interaction between expert overall rating and S&P rating, corresponds to an increase in sales by 18% (IRR 1.18).

Figure 3.6, generated using the sjPlot package [145] in R, shows the interaction between price, expert overall rating and S&P rating. We observe that an increase in expert S&P rating corresponds to marginally higher sales for higher values of overall rating (the rightmost plot of Figure 3.6). Moreover, across all three plots, we observe that effect of S&P rating on sales is higher for lower values of price, and this effect gradually diminishes as the price increases. After a certain price point, changes in the S&P rating have little

impact on sales.



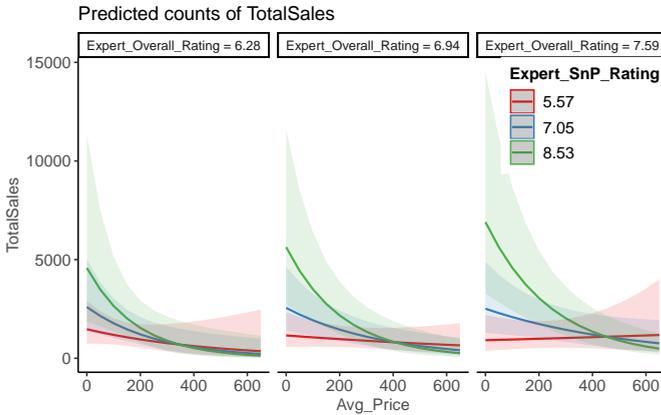
**Figure 3.5:** IRR from the Mixed effect Negative Binomial Model (\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.1$ )

This finding highlights the crucial relationship between price and S&P of a device. From our bi-variate correlation analysis (Table 3.5) we observe a positive correlation between the price and expert S&P rating for IP cameras and smart watches. The model results reveals the context of these positive correlations and suggests that an increase in S&P rating, while possibly corresponding to a marginally higher price, also corresponds to higher sales. Moreover, as observed via the interaction effects, the amplifying effect of S&P on sales is higher for lower prices, and gradually decreases as price increases. A crucial takeaway from this is that at the same price level, IoT devices with better S&P perform better in the market in terms of sales. In the next section, we evaluate whether this better performance can be attributed to the presence of S&P signals.

With regard to the other explanatory variables, we see that the highest IRR (2.18) is for total reviews on Winkel. An increase in the number of reviews on Winkel by one standard deviation corresponds to a 118% increase in sales. In contrast, an increase in the number of reviews on Amazon by one standard deviation corresponds to only a 29% increase. This aligns with the results from our bi-variate correlation analysis: the number of reviews on Winkel is better proxy for the Winkel sales data than the number of reviews on Amazon.

With respect to user ratings, we observe no significant effect of the Amazon user rating on total sales, but the IRR for the Winkel user rating is 1.42. This indicates that while a change in the user rating on Amazon does not correspond to any change in the sales, an increase in the Winkel user rating by one standard deviation corresponds to a 42% increase in sales. This contradicts with our bivariate correlation analysis (Table 3.6), which shows positive correlation for both Amazon and Winkel user ratings, suggesting that the linear relationship between Amazon user rating and sales is possibly influenced by other confounding variables. Moreover, the increase in user rating on Winkel corresponding to

an increase in sales underscores the role of consumer metrics in influencing the purchase decisions of consumers. In an online store, such mechanisms play a crucial role in signalling device quality and stimulating trust [191]. With respect to the random effect of the device type, we observe a variance of 0.234. This suggests that only a minimal amount of additional variability in sales can be explained by the IoT device type.



**Figure 3.6:** Interaction effects of Price, expert S&P rating and expert Overall rating on Total sales

### 3.5. RELATIONSHIP BETWEEN EXPERT RATINGS AND INFORMATION FROM PRODUCT PAGES

In this section, we address our second research question assessing whether user ratings and update support duration information align with expert ratings. This allows us to evaluate the presence of S&P signals and the extent of information asymmetry in the context of the purchase decisions we analysed in Section 3.4.

#### 3.5.1. RELATIONSHIP BETWEEN EXPERT S&P RATINGS AND AVERAGE USER RATING

We first analyse the correlation between the expert S&P rating – which also includes the expert update rating – and the user ratings on Amazon and Winkel. Table 3.8 provides an overview of the Spearman correlations between the expert ratings and the Amazon and Winkel user ratings. Since our primary aim is to analyse if these ratings act as signals for S&P, we do not analyse the relationship with the expert overall rating. However for the interested reader, we present the correlation between user rating and overall ratings in the Appendix B.4.

At the aggregate level, across all the device types, we observe a low negative correlation between the expert update rating and the user rating on both Amazon and Winkel. Moreover, at the individual device level, although the expert update rating for smart speakers shows a low positive correlation with both the user ratings, the update ratings of both IP cameras and smart printers show a negative correlation with the Winkel user rating. The

negative correlations could be a reflection of consumers' negative experiences with the update process [94] or a consequence of consumers not being able to account for good update practices into the user rating due to lack of visibility or information. Although invisibility is considered a desirable feature in security design [29], it can also lead to consumers not being able to account for security in their rating of the device.

**Table 3.8:** Correlation between CB S&P and update ratings and user ratings on Amazon and Winkel

Device type	Spearman correlation			
	Expert update rating vs		Expert S&P rating vs	
	Amazon user rating	Winkel user rating	Amazon user rating	Winkel user rating
IP Cameras	-0.007	-0.215**	0.012	-0.104
Smart printers	-0.134	-0.376***	0.141	0.416***
Smart speakers	0.380**	0.324*	-0.118	-0.323*
Smart watches	-0.005	0.098	0.204**	0.008
Aggregate	-0.152***	-0.23***	-0.011	-0.077

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

Regarding the expert S&P rating, we find that, at the aggregate level, across all device types, there are no significant correlations between the expert S&P rating and the user ratings on both Amazon and Winkel. At the individual device level, we observe a modest positive correlation between the expert S&P rating for smart watches and smart printers and the user ratings on Amazon and Winkel respectively. In contrast, the S&P rating of smart speakers shows a moderate negative correlation with the Winkel user rating. The negative correlation, taken together with the couple of low positive correlations shows that user rating cannot serve as a reliable indicator of the S&P posture of an IoT device. Rather, the user ratings likely reflect the overall consumer experience rather than an evaluation of the S&P posture of the device. Moreover, the differences in the correlations between the expert ratings and the user ratings across the device types is inline with earlier work that found that percentage of references to S&P varied across reviews for different IoT device types [225].

### 3.5.2. RELATIONSHIP BETWEEN UPDATE RATINGS AND UPDATE SUPPORT INFORMATION

Next, we turn to analyzing the provisioning of the update support information and the expert update ratings of CB, to better understand if the update statuses are reflective of general update practices of manufacturers. If they are, they can serve as signals for S&P and help consumers make more informed purchase decisions with respect to IoT S&P. Table 3.9 shows the mean values of expert update rating for the different update information statuses on Amazon and Winkel. We observe that on both Amazon and Winkel, devices with valid update support information have a higher update rating. To check if the differences between the two groups on Amazon (Invalid and Valid), and the three groups on Winkel (Invalid, Partly valid and Valid) are significant, we use two different tests. For Amazon, we use the Mann-Whitney U test and for Winkel, we use Kruskal-Wallis test with an additional post-hoc Dunn's test if the results are significant.

**Table 3.9:** Mean values of update rating for different values of update status on Amazon and Winkel

Update Status	CB Update rating
<b>Amazon</b>	
Invalid (397)	5.2
Valid (48)	5.9
<b>Winkel</b>	
Invalid (191)	5.1
Partly valid (151)	4.7
Valid (103)	6.4

Our results shows that the differences on both Amazon and Winkel, among the different groups are significant. The statistically significant higher expert update ratings for devices with valid update support duration information imply that the mere availability of update support duration information reflects manufacturers' practices regarding updates. This might then serve as a signal for security, reducing the information asymmetry for IoT devices. Note that the this signal is only available for a minor percentage of devices under consideration: only 22.4% of devices on Winkel and 11.8% of devices on Amazon had valid information. Moreover, for the 18 devices for which the information was available on both websites, the update support duration was not consistent. Nonetheless, we verify if there is a correlation between the presence of valid update support information and higher device sales in the next section.

### 3.6. RELATIONSHIP BETWEEN UPDATE INFORMATION STATUS AND SALES

In this section, we address our third research question on the relationship between provisioning update support information and sales. To do so, we extend the model and add categorical variables for update information status for Amazon and Winkel, in addition to the other explanatory variables. The results are shown in [Table B.2](#) in [Appendix ??](#). As mentioned earlier, on Amazon, the update status has two values, Invalid and Valid while on Winkel it has three values Invalid, Partly Valid and Valid. [Table 3.10](#) shows the IRR for

**Table 3.10:** IRR for Update Statuses on Amazon and Winkel

Update Status	IRR
Amazon Invalid (397)	0.66**
Amazon Valid (48)	0.87
Winkel Invalid (191)	0.79
Winkel Partly valid (151)	1.1
Winkel Valid (103)	1.69*

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

the update statuses. We find that, after controlling for the other factors in our model, the status of invalid on Amazon corresponds to a 34% decrease in sales (IRR 0.66) while a valid update status on Winkel corresponds to a 69% increase in sales. This implies that the presence of update support duration information, which serves as a reasonable signal

for the S&P of an IoT device (from [section 3.5](#)), also corresponds to a positive change in sales. Although the underlying causal mechanisms are unclear, this suggests that there are positive incentives at play that will reward sellers and manufacturers who adopt initiatives like IoT security labels [69, 111, 201]. Moreover, our results align with other research [157] indicating that the availability of update support duration information has an impact of 8% to 35% on consumers' purchase decisions.

### 3.7. DISCUSSION AND RECOMMENDATIONS

3

Before we discuss our results and offer recommendations, we want to emphasise that this is an observational study. We have analysed the relationships between the variables but this is insufficient to establish causality. We found that despite the lack of S&P signals for a majority of devices, there is an alignment between consumer preferences for S&P – as established in prior work [70, 90] – and their purchase decisions, although the underlying causal mechanisms are unclear. We also find that the relationship is moderated by the price of the device. At lower prices, devices with higher S&P ratings correspond to higher sales while this relationship diminishes as the price increases.

Previous studies on security labels [70, 90] have found that consumers are willing to pay more for S&P when provided with information about it. Our results corroborate these findings and additionally suggest that, in the absence of security labels, the willingness to pay more for S&P exists up to a certain price point, beyond which the demand for S&P decreases. This dual emphasis on both security and price highlights the nuanced nature of consumer preferences in the IoT market.

This also underscores the challenge for manufacturers to balance better S&P at affordable prices. As we observe in our study, at the aggregate level, across all device types, and also for IP cameras and smart watches, the expert S&P rating positively correlates with price, suggesting that increase in S&P corresponds to an increase in price. This highlights a dilemma for manufacturers. They can focus on S&P at the risk of increasing the price and losing out on sales, or they can offer devices at comparable price points and enjoy higher sales. Note that the model results show that a one standard deviation increase in price (112.7 EUR or 121.35 USD), is associated with a 27% decrease in sales. While it is beyond the scope of this study to determine what types of security controls are feasible to implement within 112.7 EUR, a crucial takeaway is that at the same price points, IoT devices with one standard deviation or 1.5 unit increase in S&P correspond to a 56% increase in sales. This suggests that manufacturers who manage the balancing act between S&P and price stand to gain significant benefits.

With regard to the user ratings, we find limited correlation between the user ratings and expert S&P ratings suggesting that only a minority of the reviews and associated ratings might express S&P concerns that align with expert evaluations. This indicates that most consumers may either lack the ability to discern the S&P features of a device or may not consider them significant factors in their evaluations and ratings. Alternatively, even when factored into ratings, their S&P concerns might not align with expert evaluations, possibly due to friction with S&P configurations like 2FA [70, 225].

This emphasises an important aspect that needs to be balanced with S&P – usability. As the results of our correlation with expert update ratings highlight, better update ratings correspond to lower user ratings ([Table 3.8](#)). This could be a consequence of consumers

not being able to account for good update practices into the user rating due to lack of information. It could also be a reflection of consumers' negative experiences and friction with the update process [94]. This emphasises the importance of designing S&P features with a focus on user experience and aiming to decrease such friction.

Moreover, increased transparency about the S&P features of an IoT device that informs users about the steps needed for the added S&P will enable better management of consumer expectations. As the results from our third research question show, on Winkel, devices for which complete update support information was available correspond to a 69% increase in sales. Although it requires further research to understand the underlying causal mechanisms for the boost in sales, our results, in line with other user studies [157], suggest that there might be positive incentives at play for IoT devices that contain S&P information at the purchase point. Crucially, availability of S&P information at purchase might increase the price point till which consumers are willing to pay more for better S&P. Thus, initiatives like security labels will yield benefits not only for the consumer, in terms of more informed purchase decisions, but may also for reward manufacturers and sellers with higher sales.

To that end, as our study results show, sellers play a key role in publishing relevant information on online stores. Three years after the mandate on provisioning update support duration information, we find that only 4.8% of devices analysed contain valid, albeit different update support duration on both stores. This implies that any effort aimed at improving manufacturer compliance, for security labels and the like, should also consider the various sellers who operate as intermediaries between manufacturers and consumers on e-commerce platforms. Although the onus is on the manufacturers to provide the information to the sellers and on the online store to enable provisioning of the information on the product pages, the task of updating the relevant information in the product pages still falls to the sellers. In the ongoing discussions on stakeholders in the realm of S&P IoT devices, these intermediaries are often overlooked. This oversight causes us to miss their perspective, which is important for the successful implementation of security labels. Unlike traditional brick-and-mortar stores, where these labels can be physically displayed on products with minimal intervention from store owners, online stores rely on sellers to present information as they deem appropriate. As we see with the case of update information, even if the online store has the provision to display this information, it often remains empty or has incomplete information. Therefore, it is vital to emphasise the necessity of transparency regarding the S&P features of a device, not only to the manufacturers, but also to the sellers.

### 3.7.1. LIMITATIONS

The expert ratings used in our study come from the tests conducted by the network of European technical test labs commissioned by various consumer associations in Europe. Although other tests using different constructs of S&P might arrive at different ratings, we believe these trusted ratings serve as a good indicator of the overall S&P posture of the device. A limitation of the testing of consumer associations is that they maximize for consumer benefit and therefore focus on devices that are popular in the market. While this introduced some selection bias in our data, the observed S&P ratings were diverse enough to meaningfully answer our research question.

Further, although our sales data is from a single e-commerce store in the Netherlands, we find that it is well aligned with the national market study data (refer subsection 3.3.2). Moreover, while any study based on a single country raises questions on generalizability, prior work estimating the popularity of IoT devices across geographies found that across all regions, 100 vendors account for nearly 90% of all IoT devices [132]. This suggests that while there might be minor differences in device popularity across regions, there is an overarching uniformity in vendor distribution which makes the implications from our findings applicable to e-commerce stores across geographies. Moreover, although including user reviews would have led to a richer analysis, we excluded them since prior work found that only 9.8% of user reviews of IoT devices contain references to S&P [225].

As we observe in our study, there is variations in the results among these four device types, which might limit their relevance to other IoT device types. We account for these variations in the model by treating the device types as random effects while studying the fixed effects of the other factors. Moreover, although the factors we included were based on prior literature, we acknowledge that there might be other confounding variables that influence consumer purchase decisions which we have not considered. Future work can evaluate the influence of other factors and deepen our understanding of the underlying causal mechanisms at play.

### 3.7.2. RESEARCH ETHICS

We got approval for the study from the Ethics Review Board in our institution. We assessed our methodology against the ethical guidelines outlined in the Menlo Report for ethical practices in computing studies [119]. We only scraped publicly available information from Amazon and Winkel, and did not collect the associated usernames. Moreover, to ease the load on the servers, rather than use web crawlers, we fed specific pages to the scraping scripts. We also added delays and distributed the scraping over a longer duration to minimise server strain.

## 3.8. CONCLUSION

In this work, we analysed consumer purchase decisions and signals for S&P in the market context within which these decisions are taken. Our results showed that despite lack of information about S&P for a majority of the devices, a one standard deviation increase in the S&P rating of a device is associated with a 56% increase in sales. However, the relationship is moderated by the price of the device. The effect is stronger at lower prices and decreases corresponding to an increase in price. In addition, we find that availability of complete update support information on online stores reflects the update practices of manufacturers and can therefore act as a signal for S&P. Further, we also find that on one of the online stores, devices with complete update support information correspond to a 69% increase in sales. This suggests that there are positive incentives at play that will reward manufacturers and sellers who adopt S&P transparency initiatives like security labels with higher sales. Our results also highlight the crucial role of sellers in ensuring the success of such initiatives since they are responsible to update the relevant information on online stores.

# 4

## LONGITUDINAL EVOLUTION OF IoT S&P

*Despite significant efforts to enhance the security and privacy (S&P) of IoT devices in recent years, concerns about their S&P remain critical. A compelling question persists: are newer IoT devices genuinely better equipped with S&P features than their predecessors? In this study, we address this gap by leveraging longitudinal S&P testing and rating data from a European Consumer Association known for conducting rigorous, expert-driven assessments of IoT S&P. We focus on three widely used and mature IoT device types – IP cameras, smart printers and smart speakers – and analyse their S&P trends over the past decade. Our findings reveal that overall, S&P ratings have remained relatively stable over the years, with some fluctuations but limited improvement. Additionally, our analysis identifies a surprising trend. There is widespread inconsistent deployment of S&P features across successive device models by the same manufacturer, suggesting a lack of systematic improvement. Results from a Beta Regression Model emphasize the pivotal role of manufacturer-level factors, such as industry experience and headquarters location, in shaping S&P outcomes. These findings underscore the need to develop comprehensive guidelines to help manufacturers operationalise and integrate S&P best practices into their development processes to promote a systematic improvement of IoT S&P.*

### 4.1. INTRODUCTION

Historically, numerous incidents have shaped the narrative of IoT devices having poor S&P, from the infamous Mirai botnet attacks [23] to a class-action lawsuit against Amazon due to multiple instances of Ring cameras being hacked [91] to the alarming case of Xiaomi cameras enabling users to view random video feeds from other users [212].

As a consequence, pressures on manufacturers have mounted to improve the security of their devices. These range from naming and shaming manufacturers [174] to regulatory investigations [57] to the development of new standards and guidelines [73, 167] and, last but not least, new legislation. For example, the EU's Radio Equipment Directive [54]

has been imposing binding security requirements on an expanding array IoT product categories and upcoming Cyber Resilience Act will generalize this approach across a much wider range of products [55]. In the US, federal and state laws have also adopted IoT security requirements [101, 195].

These recent developments lead to a critical question: are newer IoT devices genuinely better equipped with S&P features than their predecessors? Singular incidents with IoT devices in the wild cannot reliably indicate trends. The assumption that newer IoT devices would have better S&P features compared to those from a decade ago while reasonable, remains largely unverified. The lack of empirical evidence is not surprising — rigorously assessing the longitudinal S&P features of a substantial number of devices across diverse market segments is a challenging and resource-intensive task, typically beyond the scope of academic research.

4

In this paper, we address this gap by collaborating with a European Consumer Association (ECA). The partner ECA belongs to a network of consumer associations in different European countries. This network collectively had nearly half a million paid subscribers. They use this funding for device testing. ECA and its peers coordinate to systematically evaluate and rate various aspects of IoT devices, including S&P features, through a network of federated labs across Europe. We obtained a unique dataset that contains S&P ratings from 213 S&P feature tests for 428 IoT devices from 23 manufacturers of three popular and mature device types: IP cameras (released between 2016 and 2024), smart printers (released between 2014 and 2023) and smart speakers (released between 2019 and 2023). These devices were selected for testing by the ECA when they began to gain popularity a few years ago. This allowed us to access longitudinal testing data spanning from the start of the tests to 2024. Moreover, IP cameras and smart printers are overrepresented among compromised IoT devices [23], making them particularly relevant for evaluating longitudinal trends in S&P features.

Using this dataset, we aim to answer the following research question: *How have the S&P ratings and features of three types of IoT devices evolved over the past years?* We first analyse the temporal trends in IoT S&P ratings over the past decade at both the device type level, so across different manufacturers, and at the level of individual manufacturers, across consecutive models of that device type. Our analysis shows that statistically, the S&P ratings of IP cameras and smart speakers show no significant temporal trend while the ratings of smart printers show a decrease over time. At the manufacturer level, we find that only 3 of 24 manufacturers show an increase in S&P ratings over time. About half of the manufacturers have stable S&P ratings while the remaining manufacturers show erratic changes over time with no clear trends.

Next, we delve deeper into the specific S&P features that underpin the ratings. We find a surprising trend: the stable ratings actually obscure changes at the feature level, meaning that some S&P features are added and some removed from one device to the next from the same manufacturer. This indicates a lack of systematic focus on S&P in the development process of IoT manufacturers. To understand the underlying factors that influence the S&P features, we construct a Beta Regression Model. The model shows that greater experience of the manufacturer both in terms of number of models per IoT device type and the age of the manufacturer increases the odds of having a higher proportion of positive S&P features. It also shows that manufacturers with headquarters

in the US have higher odds of having positive S&P features. Overall, we find little evidence for overall improvements in S&P of IoT devices in the three device types, though with some surprising patterns at the level of specific features. In sum we make the following contributions:

- We present the first paper to systematically evaluate S&P features of IoT devices longitudinally. We analyse trends in deployment of S&P features across three IoT device types and across different models from the same manufacturer.
- We uncover a surprising trend in IoT S&P feature deployment. Basic features, such as software update support, are sometimes removed after being included, only to be reintroduced later. This pattern suggests that the S&P features of IoT devices may reflect the S&P of the underlying SDKs used or the differing S&P priorities among various product development teams, rather than a cohesive, overarching manufacturer policies.
- We develop a model to identify the factors associated with the presence of a higher number of robust S&P features in an IoT device. Our findings indicate that manufacturer characteristics have a stronger influence on the presence of S&P features than device-level factors like price.

The rest of the paper is structured as follows: [section 4.2](#) reviews related work, [section 4.3](#) provides an overview of the testing process at the ECA and [section 4.4](#) contains the dataset description. In [section 4.5](#) we present the analysis of trends in S&P ratings and in [section 4.6](#), the trends in S&P features. In [section 4.7](#), we construct a model to explore factors influencing IoT S&P. In [section 4.8](#) we discuss our findings, offers recommendations and address the limitations, and conclude the paper in [section 4.9](#).

## 4.2. RELATED WORK

To the best of our knowledge, there are no prior papers that have investigated the trends in S&P across IoT device types or the evolution of S&P features on IoT device models from the same manufacturer. The closest is a study by Dong et al. [67] who analysed the Transport Layer Security (TLS) of IoT devices from the same vendor and found only a small fraction of TLS fingerprints are reused by vendors. Other papers have systematically evaluated the S&P of a wide range of IoT devices among which are some devices from the same manufacturer. For instance, Loi et al. [143] conducted independent tests of twenty IoT devices to develop a comprehensive security overview. Among these twenty, there are two manufacturers with two devices each and one manufacturer with three devices. Within this limited dataset, the security test results look quite similar. In contrast, a large scale analysis of IoT devices on home networks revealed that within a set of devices made by the same manufacturer, the percentage of devices with weak default credentials has a wide range from 1.4% to 97.1% [132].

### 4.2.1. IOT VENDOR STUDIES

There have also been some prior work studying security practices, like patching, of IoT manufacturers and vendors. Nakajima et al. [162], in their pilot study, found that five out

of the six vendors released patches on time. Pérez et al. [183] find similar results with a larger pool of 104 vendors and highlight that IoT-centric vendors release more patches on time than non-IoT-centric vendors. Further, they find no significant relationship between vendor size and patch availability suggesting an absence of economies of scale.

#### 4.2.2. IOT SECURITY BEST PRACTICES

There are many studies that aim to assist manufacturers in improving the S&P of their IoT devices. For instance, the US National Institute of Standards and Technologies (NIST) has outlined foundational security practices of manufacturers with the user at the centre [77] while other solutions are centred around the IoT product life cycle [233]. Bellman and van Oorschot [30] note that 70% of the best practices they analysed in their literature review related to early IoT device life cycle, highlighting the crucial role of manufacturers in addressing IoT S&P.

A survey conducted by Akiyama et al. [15] among IoT professionals highlighted the complexity of the IoT software supply chain as one of the challenges in managing the IoT software components. Morgner and Benenson [155], through a case study on IoT standardisation, notes that large investments in security are not prioritised due to a lack of consumer demand.

#### 4.2.3. SECURITY PRACTICES IN ORGANISATIONS

Other studies focus on understanding the S&P development processes within software development in general and not specific to IoT manufacturers. Through interviews with software developers, Assal and Chiasson [25] discovered that best practices in literature often overlook factors involving the team's operational strategies, company culture and security knowledge. Xiao et al. [231] argues that software development companies as social systems and the social factors like policies and cultures influence adoption of secure development tools. Similarly, Arizon-Peretz et al. [24] note that factors like organisational security climate, individualism vs collectivism, security self efficacy and proactive security behaviour influence the adoption of security by design principles.

In sum, our literature review underscores the lack of studies on evolution of S&P in IoT devices or on the evolution of the S&P development processes of manufacturers.

### 4.3. ECA TESTING METHODOLOGY

In this section, we provide an overview of the testing methodology followed by our partner, a prominent consumer association in Europe, who wished to remain anonymous. The partner European Consumer Association (ECA) has around 420,000 paid members as of January 2025. The members get exclusive access to test results of different products including IoT devices. Since these test results are standardized across all European consumer associations, we cannot publicly share the specifics, as manufacturers might exploit this information to achieve a higher S&P rating instead of focusing on holistic improvements. However, we outline the process followed by the ECA, from including a new device type in testing to publishing the results on the website to highlight the rigour in the testing process.

For each device type tested, the ECA assembles a group of experts from various con-

sumer associations across Europe. These experts collaboratively identify key feature/test categories like those mentioned in Table 4.1 and develop a list of fine grained features or tests specific to each device type and category. Each fine-grained feature is designed to be tested with a binary yes/no result for clarity and consistency. For example, under the category of ‘Password Policy’, one feature might be having a unique default password, which would be mapped to a specific test question ‘Does the device have a unique default password?’ This structured approach ensures that each overarching category is rigorously evaluated through multiple precise tests. The number and type of tests vary by category and device type; for instance, while password policy checks are standard across all devices, unique assessments like data security for smart speakers target specific usage contexts. This does not imply that data security is overlooked for IP cameras or smart printers, but rather that ECA prioritizes features most relevant to consumers for each device type. This feature identification and test development is an iterative process and the result is a comprehensive list of tests for each device type. Once the list of tests is identified, the experts also assign a weight to each test based on its relative importance. For example, a unique default password carries more weight than merely supporting ASCII characters. Once the team of experts agree on the tests and weightage, it is passed on to the network of federated labs where the testing is done.

**Table 4.1:** List of feature/test categories for each device type

IP Cameras	Smart Printers	Smart Speakers
Password Policy	Setup	Data Security
Standard Installation	Access Controls	Decommissioning
Android App	Password Policy	Password Policy
iOS App	Updates	Network Security
Updates	Permissions	Update Policy
Known Vulnerabilities	Encryption	iOS App
	Authentication	Android App
	Known Vulnerabilities	Privacy Policy
	Decommissioning	

This structured approach ensures objective, consistent testing, limiting subjective interpretation by lab technicians. While other testing methodologies using different S&P constructs might yield different results, we maintain that the rigorous, expert-developed ECA tests eliminate subjectivity and thoroughly assess critical S&P aspects and provide a reliable evaluation of a device’s overall S&P posture. Hence, we use these ratings, and the underlying test results in conjunction with the conversion scale – to determine which test outcome gets a higher rating and is therefore better for S&P – in our analysis.

#### 4.4. DATASET DESCRIPTION

The ratings collected are for three IoT device types – IP cameras, smart printers and smart speakers. In addition to the ratings, we also obtained the binary results from the underlying tests. Since we aimed to analyse manufacturer-level changes in S&P, we use data only from manufacturers with three or more devices. This approach ensures consistency in the data used for both trend and manufacturer analyses, and allows us to meaningfully analyse the trends within individual manufacturers and also the broader trends across manufacturers. In addition, we collected meta data like the release dates

Device Type	# Tests	# Mfgs.	Avg. # Devices	Price Range	Release Date Range	Mfg. HQs	Founding Year Range	Mfg. Size Range
IP Cameras	79	14	8.8	\$29.12 – \$832.00	Jan 2016 - Jan 2024	China, Sweden, Taiwan, USA	1840 - 2016	108 - 182,502
Smart Printers	60	6	53	\$46.70 – \$861.10	Jan 2014 - Nov 2023	Japan, USA	1906 - 1991	9000 - 184,034
Smart Speakers	74	3	4.3	\$29.50 – \$665.60	Aug 2019 - Jan 2023	The Netherlands, USA	1943 - 2002	1,844 - 1,525,000

**Table 4.2:** Overview of the dataset

4

to analyse the temporal dimension of S&P ratings. In some cases, the release dates were available in the ECA dataset, for others used other sources such as manufacturer websites, consumer product launch blogs, date first available on Amazon and so on. Moreover, since prior research has highlighted the importance of organisational factors in security [127, 138], we also collected four variables at the manufacturer level: the headquarters location, founding year and the employee size. These allow us to understand the relative influence of each of these factors in the S&P of IoT devices.

Table 4.2 gives an overview of the dataset for the three device types. The data set contains 14 IP camera manufacturers, 6 smart printer manufacturers and 3 smart speaker manufacturers. The number of devices per manufacturer varies both between manufacturers within each device type and across different device types. On average each IP camera manufacturer has 8.86 devices, with a median of 6 devices. In contrast, manufacturers of smart printers have an average and median of 53 devices, whereas smart speaker manufacturers have an average and median of only 4 devices. This stark difference reflects underlying market dynamics. The smart printer sector features a broader range of products from many manufacturers, resulting in a higher number of devices per manufacturer. In contrast, the smart speaker market is more concentrated, dominated by a few major players who offer limited product lines with mostly incremental updates.

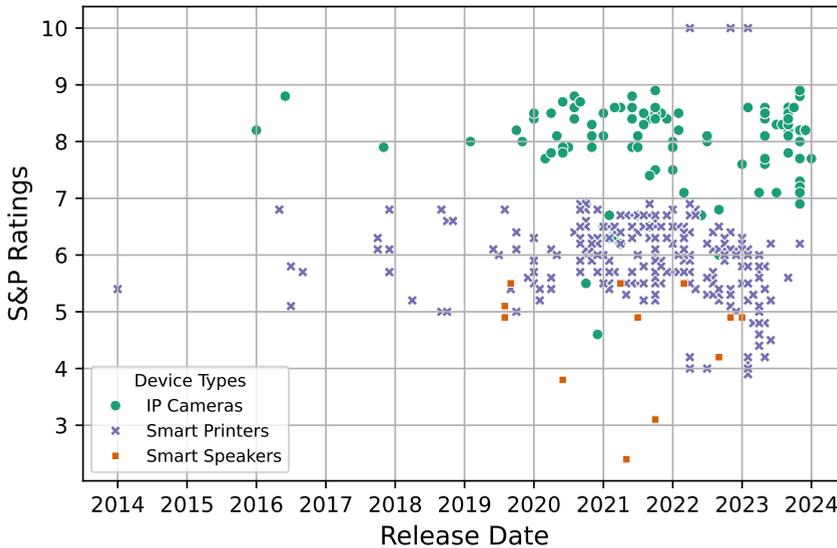
The release date ranges differ per device type, with smart printers having the longest range of almost a decade, IP cameras having a range of eight years and smart speakers having the smallest range of about three and a half years. This also reflects the time periods when the device types saw an increase in popularity. With regards to the headquarters (HQ), almost half of the manufacturers (11/23) have HQ in the US, six in China, three in Japan and one each in the Netherlands, Sweden and Taiwan. Both the oldest (established in 1840) and the newest (established in 2016) manufacturers are IP camera manufacturers. The smallest manufacturer in terms of employees is an IP camera manufacturer with a mere 108 employees while the largest manufacturer is a smart speaker manufacturer with around 1.5 million employees.

## 4.5. TRENDS IN IOT SECURITY AND PRIVACY RATINGS

In this section, we address our first research question regarding the evolution of S&P ratings of three IoT device types over the years to understand the underlying trends. We examine trends at two levels: at the device type level across all manufacturers and also at the manufacturer level across different IoT device models from the same manufacturer.

### 4.5.1. EVOLUTION OF S&P RATINGS AT THE DEVICE TYPE LEVEL

Figure 4.1 shows the distribution of S&P ratings of the three IoT device types over the years. The test results are available from 2014 for smart printers, from 2016 for IP cameras, and from mid-2019 for smart speakers. These dates reflect when each device type gained popularity, prompting the European Consumer Association to include them in their testing. Moreover, within the popular device types, the ECA chooses to test only the specific device models that are popular in the European market in order to maximise the effectiveness of the results for their members. The higher popularity of IoT devices tested by consumer associations has also been empirically verified in an earlier study [223].



**Figure 4.1:** Temporal Evolution of S&P Ratings of IoT devices

We observe that the S&P ratings for each device type have different ranges. Smart speakers have the narrowest range and the lowest S&P ratings, ranging from 2.4 to 5.5, with an average of 4.39. Smart printers have a broader range of ratings from 3.9 to 10, with a higher average of 5.98. The S&P ratings of IP cameras range from 4.6 to 8.9, with an average of 7.9. This suggests that S&P profiles vary across IoT device types, highlighting the need to consider each type individually rather than treating grouping all IoT devices in a single category.

In order to systematically evaluate the trends in the S&P ratings of these devices, we

conducted a statistical trend analysis using Mann-Kendall test, a non-parametric test that is commonly used to analyse time series. We used the `pyMannKendall` library in Python [103] to run the test for each device type separately. Our analysis showed no trend in the S&P ratings of IP cameras and smart speakers, and a decreasing trend for smart printers. This is inline with industry trends that show that printers were 68% more likely to be a source of a threat or breach in 2023 when compared to 2016 [207].

#### 4.5.2. EVOLUTION OF S&P RATINGS AT THE MANUFACTURER LEVEL

Next, we study the temporal trends in the S&P ratings of different IoT devices from the same manufacturer. We analyse the ratings of each device type separately to understand if there are differences in manufacturer behaviour across the device types.

## 4

### IP CAMERAS

The IP cameras in our dataset come from 14 different manufacturers. Each manufacturer has between 3 and 24 IP cameras. The average S&P rating for the 14 manufacturers ranges from 6.15 to 8.65. Figure 4.2 shows the evolution of S&P ratings and price of IP cameras over time. We observe three distinct temporal patterns. First, for nine manufacturers, IP\_M1 to IP\_M9, the S&P ratings have remained stable throughout the analysis period. Notably, all of these manufacturers started with high initial S&P ratings that have consistently stayed in the high ranges. The average S&P rating for the manufacturers in this group – 8.3 – is the highest of all three patterns. Second, three manufacturers, IP\_M10, IP\_M11 and IP\_M12, show some fluctuations in the S&P ratings over time, with an average rating of 7.4. In the third pattern two manufacturers, IP\_M13 and IP\_M14, both started with a low S&P rating but have shown consistent increase over the observation period. Manufacturer IP\_M13 has an average S&P rating of 6.2 and IP\_M14 has an average S&P rating of 6.9. Overall, we find no decreasing trends in the S&P ratings of these manufacturers. Most manufacturers have maintained a stable rating over the years, a few show some fluctuations while a couple exhibit a steady increase in their S&P rating.

We observe fluctuations in the prices of IP cameras released by each manufacturer over the years, possibly due different market segments being targeted with the differential pricing. Consistent with earlier studies, we also find a statistically significant mild positive correlation between the price of IP cameras and the S&P rating [225].

### SMART PRINTERS

There are six smart printers manufacturers in our dataset each with between 6 and 107 printers. We observe two temporal trends in the evolution of S&P ratings over time (Figure 4.3). First, PR\_M1 and PR\_M2 have relatively stable ratings, while the ratings of the other manufacturers fluctuate. Notably, we do not observe any manufacturer with a consistent increase in S&P ratings. Moreover, unlike IP cameras, smart printer manufacturers with stable ratings do not have high S&P ratings. The average S&P ratings of PR\_M1 and PR\_M2, with a relatively stable S&P rating, is 6.3, while the average rating of the other four manufacturers is 5.9. PR\_M1 has the lowest average S&P rating of 5.6, while PR\_M5 has the highest with 6.3. We do not observe any statistically significant correlation between the S&P ratings and prices of smart printers.

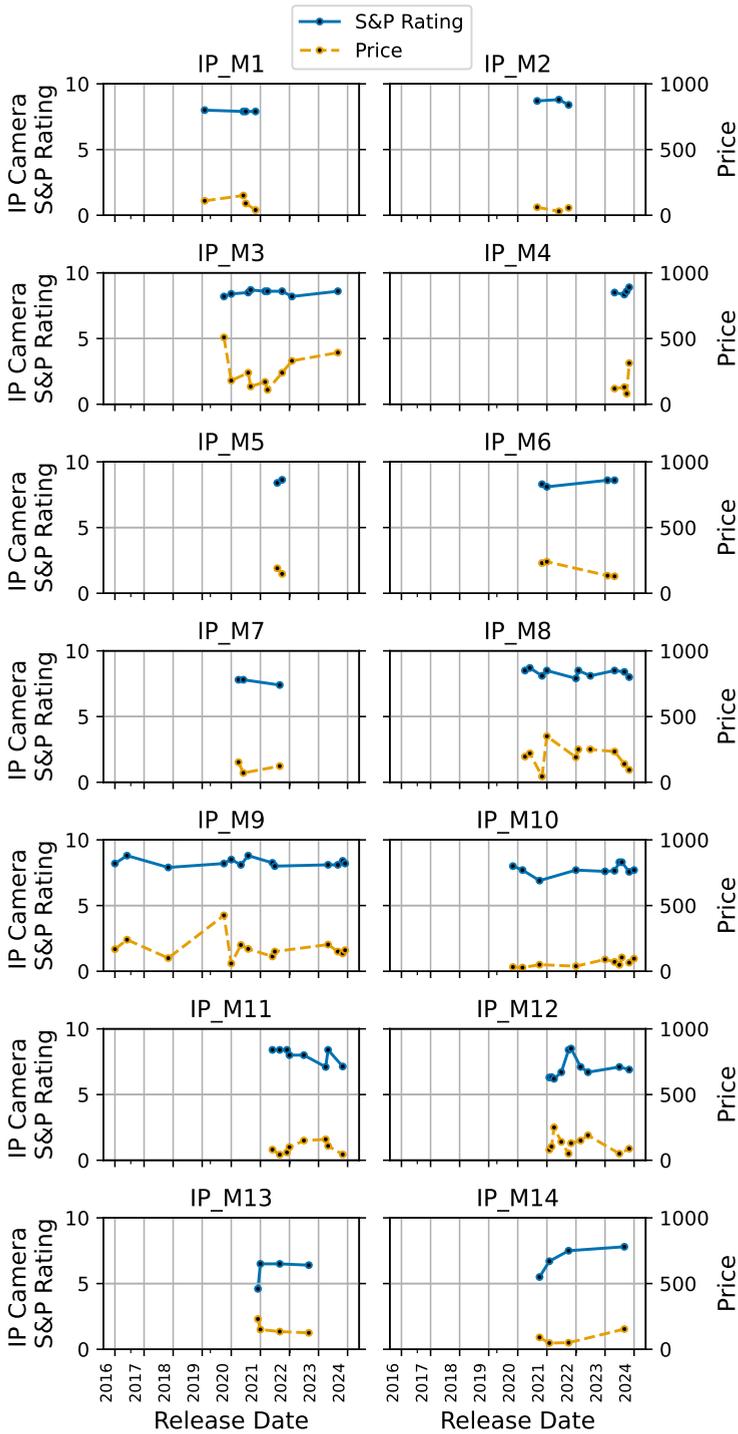
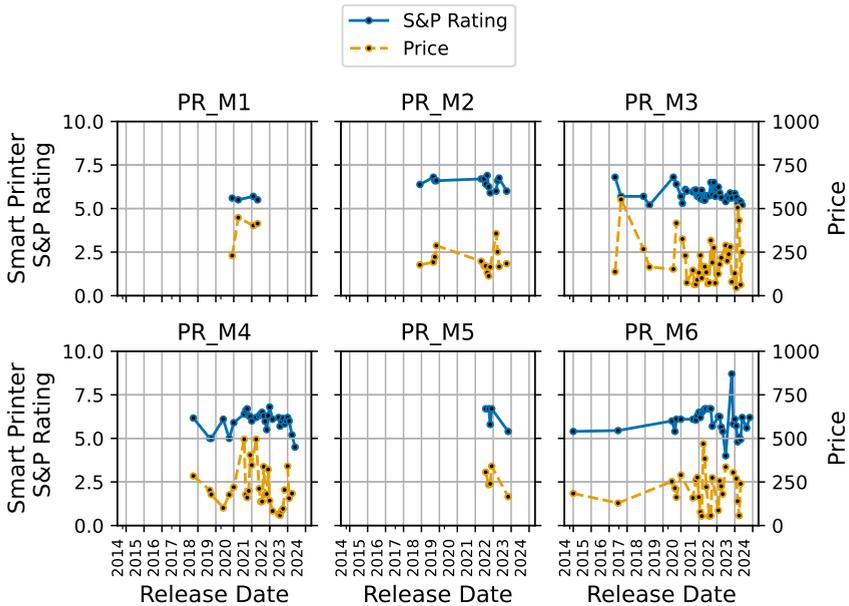


Figure 4.2: Evolution of S&P ratings and Price of IP cameras per Manufacturer



**Figure 4.3:** Evolution of S&P Ratings and Price of Smart Printers per Manufacturer

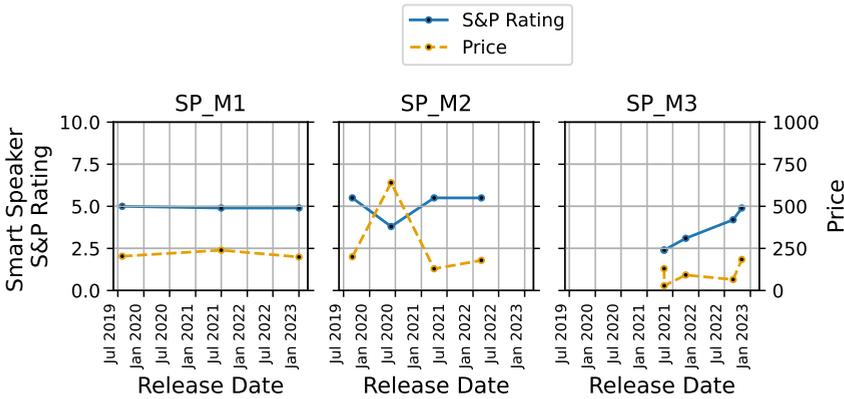
### SMART SPEAKERS

There are only three smart speaker manufacturers in our dataset, two with four devices and SP\_M3 with five devices. The S&P ratings of the four devices from manufacturers SP\_M1 and SP\_M2 are almost identical with only a maximum of a 1.5 point variation. The five devices from SP\_M3 show a higher variance. The average S&P ratings for the three manufacturers are 4.9, 5 and 3.4 respectively.

We find that each smart speaker manufacturer has a distinct temporal pattern (Figure 4.4). Manufacturer SP\_M1 maintains a stable S&P rating over time, with an average of 4.9. In contrast, manufacturer SP\_M2, with an average of 5.1, shows a steep decline in the ratings of one of its devices. Manufacturer SP\_M3 shows an increasing trend in S&P ratings, with an average of 3.4. The number of smart speakers in our dataset is insufficient to conduct statistical tests. However, visually, we observe a negative correlation between S&P ratings and price, consistent with earlier studies [225], particularly evident for manufacturer SP\_M2.

## 4.6. TRENDS IN IOT SECURITY AND PRIVACY FEATURES

In this section, we analyse the evolution of S&P features captured by the tests. Similar to the ratings, we analyse both at the device type level and at the manufacturer level.



**Figure 4.4:** Evolution of S&P ratings and Price of Smart Speakers per Manufacturer

#### 4.6.1. EVOLUTION OF S&P FEATURES AT THE DEVICE TYPE LEVEL

At the device type level, we analyse the evolution of the S&P features to get a broader perspective of where the industry is moving in terms of IoT S&P.

##### IP CAMERAS

We find that IP Cameras have consistently had good S&P features in certain areas especially in Data Privacy and Permissions. For instance, all the cameras in our study have a thorough account deletion feature and do not leak any user data on account deletion. The apps on both Android and iOS allow for opting out of data collection, and for usage of main app features even when not all requested permissions are provided. However, we also find some room for improvement. For example, secure remote connections are still not commonplace, and apps often fail to explain the consequences of declining the privacy policy or denying permission requests. Additionally, support for some features has changed over the years. For instance, while all IP cameras since 2016 have update support, between 2017 and 2022 there was a decrease in the number of IP cameras providing automatic security updates. Moreover, since 2017, lesser number of devices provide a manufacturer statement about the duration of update support. On a more positive note, there has been a marginal increase in the number of devices supporting offline functionality since 2019.

##### SMART PRINTERS

We find that only a few S&P features in smart printers have been consistently steady since 2014. This includes offline functionality, encrypted network communication, and online user access control. Some features have not improved over the years; for instance, communication is still only encrypted on the device and not on the mobile apps. Users lack options for advanced content encryption, and secure connection protocols are generally not used. Moreover, secure communication support in most mobile apps has declined since 2021.

Nevertheless, we observe an improvement in certain S&P features over time. For

instance, since 2019, an increasing number of printers offer automatic updates via the web portal. Regarding data privacy and permissions, we observe variations in the proportion of devices with privacy-friendly features. Between 2014 and 2016, all printers had personalised marketing disabled by default. This dropped to 20% by 2018, rose again by 2020, and has since declined – highlighting the fluctuating state of S&P features in smart printers.

### SMART SPEAKERS

The number of smart speakers in our dataset is much smaller than that of IP cameras and smart printers, a possible artifact of the high market concentration for smart speakers. We find that smart speakers have steady and robust S&P features in many areas, such as updates, account deletion, and factory reset — every smart speaker in our dataset supports updates and has automatic updates enabled. However, none of the speakers include a manufacturer statement about updates. Similarly, there is room for improvement in privacy and permissions, as none provide information to consumers on the consequences of denying access. Offline functionality was introduced in 2021-22, with about half of the devices supporting it, but this has since dropped to none.

4

#### 4.6.2. EVOLUTION OF S&P FEATURES AT THE MANUFACTURER LEVEL

Next, we study the changes in the S&P features between consecutive IoT device models from the same manufacturer. We define a ‘change’ as any feature variation between two consequent devices released by the same manufacturer. For instance, considering manufacturer SP\_M1 who released smart speakers in July 2019, July 2021 and January 2023, we categorise as change, any feature change between the 2019 and 2021 models or between the 2021 and 2023 models. Since each test captures a unique aspect of a S&P feature, we use the term tests and features interchangeably based on the context for ease of reading.

We further classify each change as increasing, decreasing or mixed based on how it impacts the S&P rating. The ECA uses a conversion scale to translate the test results to S&P ratings. Based on this scale, we determine whether a feature change leads to an increase or decrease in the S&P rating. For instance, if a manufacturer has six devices and the first two devices do not support automatic updates but the last four do, we would categorise this as an increase in S&P since ECA’s conversion scale gives a higher S&P rating to devices that support automatic updates. Conversely, if the first two devices support automatic updates but the last four do not, we would categorise this as a decrease in S&P since the lack of automatic updates in the latter devices would lower the rating. However, not all feature changes are consistent and we categorise these as mixed. For example, if the first two devices support automatic updates and the next two devices do not support while the last two do, we categorise this as mixed. Additionally, we also have a stable category to denote features that have remained stable over the duration of analysis. These classifications helps analyse a manufacturer’s commitment to S&P which would be reflected in the trends of feature changes across their devices over time.

Further, for each device type, we analyse the trends in specific features (eg: update support) that show a high degrees of changes across all manufacturers. This gives insights on specific features that are challenging for manufacturers to implement consistently

across their devices.

IP CAMERAS

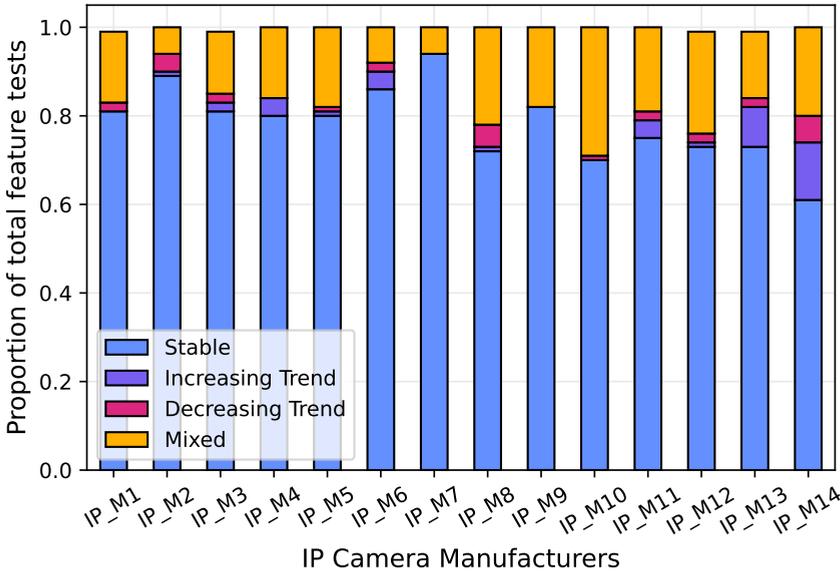


Figure 4.5: Proportion of feature changes per IP camera manufacturer in each temporal trend

In total there are 79 features that are tested for IP cameras and 60% of the features are stable across all devices from the 14 IP camera manufacturers. However, we also find all manufacturers have at least a few features that change between consecutive device models, the lowest being IP\_M7 with 6.3% of feature changes and the highest IP\_M14 with 39.2% changes between device models. We find a vast majority of these feature changes per manufacturer – ranging from half to all – have a mixed trend: they fluctuate between leading to an increase or decrease in ratings. IP\_M13 has the highest number of changes (7 of 21) that led to an increase in rating while IP\_M2 has the highest proportion of changes leading to a decrease (3 of 9). Three manufacturers do not have any changes that led to a decrease in S&P rating.

Figure 4.5 shows a stacked bar plot of the changes in for the four categories across the 14 manufacturers. Despite the observed positive correlation between the S&P rating and price at the device type level, we do not find any strong indication of price-based deployment of S&P features. Next, we analyse the features that have a high number of changes across all manufacturers – automatic update support, open ports on a device, unauthenticated password reset, offline functionality, encrypted cloud communication and app permissions.

**Automatic Update Support** Automatic update support is crucial to ensure that an IP camera stays up-to-date with security patches. However, only 5 of 14 manufacturers

have information on automatic update support for all their devices. For the remaining 9 manufacturers, the lab was not able to find the information for at least one of the devices, and in the case of IP\_M8 and IP\_M11, for about half of their devices. This highlights the difficulty of finding information related to automatic updates, and update support in general. When trained professionals struggle to find the information, it would be even harder for consumers. This aligns with another study that notes that information related to update support is not easily available even in trusted sources like user manuals and manufacturer websites [221].

Further, only two manufacturers, IP\_M6 and IP\_M9 offer automatic update support on all of their devices. In contrast, two other manufacturers, IP\_M11 and IP\_M12 - based on available update support information - do not offer automatic update support on any of their devices. Apart from these four manufacturers, all the other 10 manufacturers have at least one device that supports automatic updates and one device that does not. This inconsistency is intriguing, because if a manufacturer has the technical capacity to implement automatic updates on one device, they could potentially extend that capability to all subsequent devices. However, we see clear evidence of such capacity building over time in only two manufacturers. IP\_M10 and IP\_M13 did not initially offer automatic updates but introduced the feature in their later models. For the other eight manufacturers, the support for automatic updates remains inconsistent across their devices, reflecting a lack of systematic implementation.

**Unnecessary Open Ports on a Device** It is a common best practice to close unnecessary ports – those not needed for network functionality – to prevent unauthorised access. We find four patterns. Five manufacturers do not have any unnecessary open ports on any of their devices. Four manufacturers show a positive trend of improvement over time: their initial devices have open ports but not the latter devices. One manufacturer, IP\_M13, has open ports on all their cameras. Lastly, four manufacturers have open ports on some of their cameras.

The first three patterns could be attributed to a manufacturer's security posture. Strong focus on security leading to a 'deny by default' stance, evolving security practices causing newer devices to have no unnecessary open ports, and a higher priority placed on ease of use reflected in all devices having open ports. The last pattern, with open ports on some but not others could indicate a lack of a strong manufacturer wide security posture.

**Unauthenticated Password Reset** Since most IP cameras do not have a UI interface, they are mostly managed through an app. While some IP cameras support password reset through the app or via sending an email to the registered email id, some other cameras do not have this support. Instead, the only option if a consumer has forgotten the password, is to reset the device to factory settings, typically done through a physical reset button on the camera. This allows anybody with physical access to the IP camera to reset the device and set a new password and is hence a vulnerability from the security perspective. Only six manufacturers do not allow such unauthenticated password reset on any their cameras, while five manufacturers have addressed it over time - their latter cameras block unauthenticated password reset. With the other three manufacturers, we again find an inconsistent pattern. IP\_M1 and IP\_M14 allow unauthenticated password reset on all

their cameras except one while IP\_M9 allows it on 9 of 24 devices.

**Offline Functionality** The primary functionality of IP cameras of four manufacturers is unavailable when the internet connection is lost, even if the loss is only temporary due to a lack of offline storage. Further analysis of the devices, based information on the internet like user manuals, shows that two manufacturers IP\_M2 and IP\_M3 support local storage via a dedicated device that has to be purchased separately while IP\_M7 and IP\_M11 only support cloud storage. In contrast, all devices from manufacturers IP\_M5 and IP\_M6 can be used even when the internet connection is lost temporarily. The other ten manufacturers have inconsistent support - some devices work through a temporary internet connection loss and some that do not. Of these, three manufacturers supported basic offline functionality in the earlier devices but not in the later devices. In contrast, only one manufacturer did not support offline functionality in the earlier devices but supports them in the later devices.

Three of the four manufacturers whose devices cannot perform the primary functionality without internet, also cannot perform basic functions if the cloud service permanently shuts down. All devices from IP\_M5 and IP\_M6 can withstand temporary loss of connectivity and two of the six devices from IP\_M6 are also robust to the cloud service shutting down. We observe both an improvement and decline in offline functionality in face of cloud service shutting down over time. IP\_M4 did not support offline functionality in the initial devices, but introduced it in the latter devices while three other manufacturers had offline functionality in the earlier devices but not in the latter devices. For three other manufacturers, some devices support primary functionalities even if the cloud service were to shut down while the others do not.

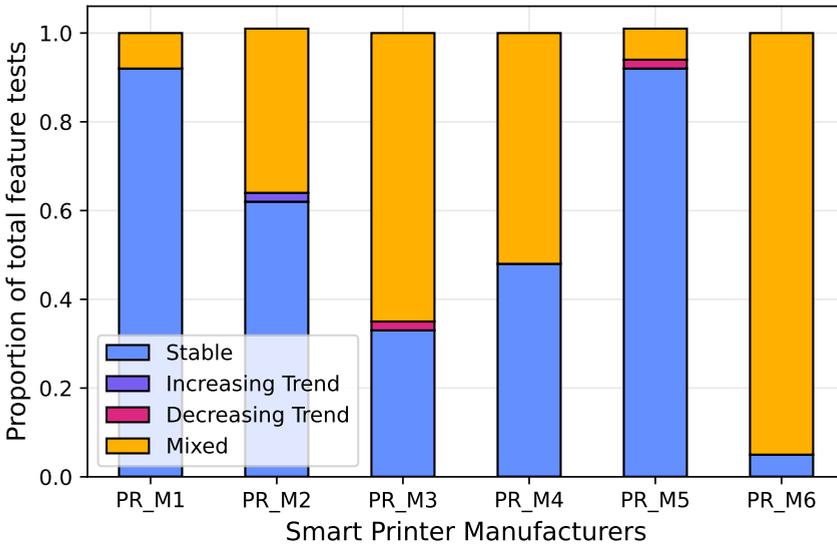
**Encrypted Cloud Communication** Most IoT devices connect to the cloud either for basic or advanced functionality. Due to a lack of user interface on most IoT devices, there are corresponding mobile apps that allow a consumer to configure and manage their IoT device, and enable remote access. The app typically connects to a cloud service for multiple functions – to act as a gateway to the device itself, to provide advanced data processing and analytics, for application hosting and so on. From a S&P perspective, it is crucial that - irrespective of the function - the communication with the cloud is secure, although the consequences of lack of secure communication might be worse in some functions than others.

In our analysis, we observe differences in the implementation of encrypted cloud communication between android and iOS apps. Half of the manufacturers offer encrypted communication on both android and iOS apps while two manufacturers support encrypted communication only on the iOS app and not on the android app. Moreover, while most manufacturers encrypt Personally Identifiable Information (PII) on both apps, three manufacturers send PII unencrypted for some devices. Of these three manufacturers, two manufacturers show a decline over time, they encrypt PII on their initial devices but not on the latter ones, while the other manufacturer shows improvement over time - their latter devices support PII encryption.

**App Permissions** We also observe differences between android and iOS apps in terms of the app permissions. While nine manufacturers do not request excessive permissions - more permissions that what is needed for device functionality - on either app, six manufacturers ask for excess permissions on the android app but not on the iOS app. Of these, two manufacturers have a positive trend: their initial devices ask for excessive permissions but the latter devices do not.

### SMART PRINTERS

Among the three device types under consideration, smart printers have the lowest proportion of features that are stable and the highest proportion that change over time. There are 60 feature tests for smart printers, and only one feature has remained consistent across all devices of all manufacturers. The remaining 59 features have changed in at least once printer from each manufacturer. Moreover, only one manufacturer (PR\_M2) has a change that led to an increase and two manufacturers (PR\_M3 and PR\_M5) have a change each that led to a decrease. All the remaining 56 feature changes show a mixed pattern alternating between leading to increase and decrease in ratings. The proportion of features that change varies across manufacturers. PR\_M6 has changes in 57 of 60 features (95%), while both PR\_M1 and PR\_M5 have changes only in five features (8.3%). [Figure 4.6](#) shows the proportion of tests of smart printers that are stable and the proportion with changes that are increasing, decreasing or mixed. Similar to IP cameras, there is no correlation between the price of the printer and the availability of the S&P features tested. We now analyse features that have a high number of changes across all manufacturers: setup, web portal access, updates and mobile app privacy.



**Figure 4.6:** Proportion of feature changes per smart printer manufacturer in each temporal trend

**Setup** With respect to the setup of the printers, we find an inconsistency between devices from the same manufacturer in the modes supported during setup. Only two manufacturers create a WiFi access point for setup across all their devices. The remaining manufacturers create a WiFi access point only for a subset of their printers. However, this highlights the trade-off between ease of use and security: most devices that allow creating WiFi access points during setup can also connect to WiFi networks without a password or without security standards like WPA2.

Using HTTP for the admin configuration exposes sensitive information to interception and tampering due to the lack of encryption and integrity checks and has more been replaced by the more secure HTTPS. However, two manufacturers still use HTTP by default for admin configuration on all of their devices and the other four manufacturers use it on a majority of their devices despite supporting HTTPS. It is puzzling that manufacturers who have the capacity to use stronger, more secure protocols like HTTPS would still use HTTP as default in a majority of their devices. On similar lines, we find that only one manufacturer supports encrypted communication with the mobile app on all of their devices, the other support it on a majority of their devices but not all devices. Similarly, only three manufacturers have up-to-date TLS certificates on all of their devices, the others do not have up-to-date certificates in around 1 to 7% of their devices.

**Web Portal Access** Although brute force attacks are a common threat to web applications, only one printer manufacturer, PR\_M4, protects against brute force attacks on the web portals of all their devices. In contrast, PR\_M6 does not support brute force protection on the web portal on any of their devices while the other four manufacturers support on some devices but not on the rest. Of these, only PR\_M2 shows an improvement over time, introducing brute force protection their later devices.

By default, the web portals of all devices from two manufacturers, PR\_M1 and PR\_M5 do not have any default password - anybody on the same network with the IP address would be able to connect to the device. In contrast, PR\_M2 has a default password for all their devices which is typically printed out of plain sight on the printer itself like on the bottom or in the battery cavity. The remaining three manufacturers have a default password for some of their devices but not for the others. Moreover, once configured, the authentication systems of PR\_M1 and PR\_M5 ask for both username and password, while the rest of the manufacturers use both only for some of the devices, the other devices ask for only a password. With regard to password reset, across all their devices, both PR\_M1 and PR\_M5 allow a password change without entering the previous, while PR\_M4 requires the older password for a password change. The other three manufacturers require the old password on some of their devices but not on the other devices. This high level of variance in the implementation of even basic S&P functionalities like setup and passwords highlights a lack of a standardised approach to S&P at the manufacturer level.

**Updates** Five of the six manufacturers support one-click updates on some printers but not on the others. For the sixth manufacturer, no update procedure related information could be found. Moreover, none of the manufacturers offer consistent support on both the web portal and mobile app for automatic firmware updates. PR\_M2 supports configuring automatic firmware updates from the web portal for a minority of devices but not from

the mobile app. PR\_M3 supports configuring automatic updates from both but only on a subset of devices. The remaining four manufacturers do not support automatic firmware update configuration via the web or app.

**Mobile App Data Privacy** The mobile apps of the printers from all of the manufacturers ask for a wide range of permissions from WiFi and fine location access to camera and cloud messaging permissions. Although some of these permissions might be not crucial for the basic functioning of the app, only the app of PR\_M1 works across all devices even when the permissions are declined. In contrast, none of the devices of PR\_M2 work if the permissions are declined. The apps of the other manufacturers work even when permissions are declined in some devices but not in the others.

With regards to the privacy policy, for none of the devices of PR\_M1, the corresponding apps explain what happens if the privacy policy is declined. For the other manufacturers, it is explained for some devices but not for the others. From a user perspective, especially if it is an inexperienced user, such explanations are useful to understand the trade offs between privacy and functionality. A lack of consistency even among devices from the same manufacturer might cause frustrations to such users.

A deeper look at the app SDK reveals that targeted advertising is turned off by default for the apps of all devices of PR\_M1. The other manufacturers have it turned off by default in some devices and turned on by default in other devices. This suggests that these decisions are probably made at the product level rather than based on a organisation wide S&P policy.

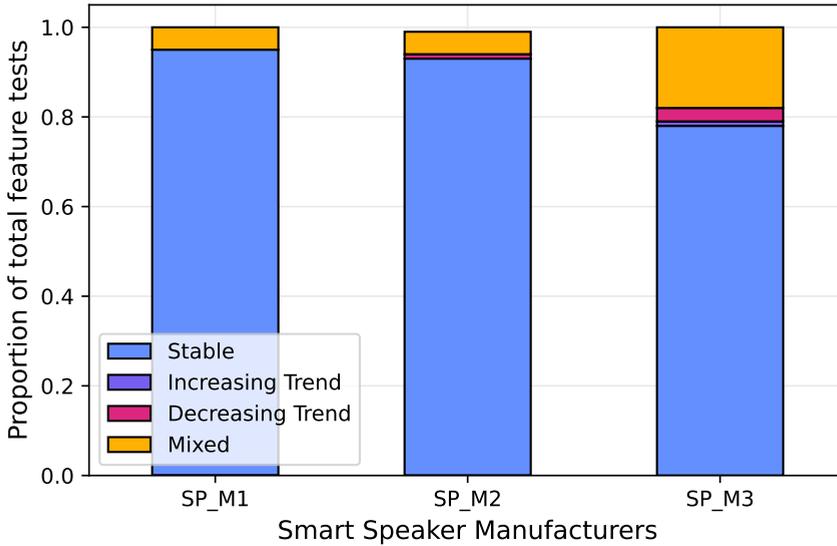
#### SMART SPEAKERS

Of the three device types under consideration, smart speakers have the lowest proportion of tests with changes. There are 74 tests for smart speakers and 58 of these (78%) remain stable across all devices and manufacturers. Moreover, a majority of the feature tests with changes have a mixed trend. SP\_M1 has a mixed trend for all changes. SP\_M3 has one feature with an increasing trend and two with a decreasing trend while SP\_M2 has one feature with a decreasing trend, the remaining feature changes show a mixed trend. [Figure 4.7](#) shows the distribution of changes in each of the trends for smart speakers. SP\_M3 has the highest percentage of features with changes – 21.6% (16 of 74) while SP\_M1 and SP\_M2 have only 4 and 5 features with changes (5.4% and 6.8% respectively).

The S&P settings of smart speakers across manufacturers are relatively homogeneous. The main point of difference is the availability of offline functionality. While none of the devices of SP\_M1 support any offline functionality, one of the four devices from SP\_M2 - the most expensive, costing three times more than the next most expensive speaker - does support offline functionality. Two out of five devices from SP\_M3 also support offline functionality, but there is no correlation with price.

## 4.7. FACTORS INFLUENCING IOT SECURITY AND PRIVACY

To understand the factors influencing S&P features of IoT devices, we constructed a Generalized Linear Model (GLM) using the `glmmTMB` package in R [41]. The dependent variable in our analysis is the proportion of positive S&P features for each device. For example, if a smart printer achieves positive results, meaning results that are better for



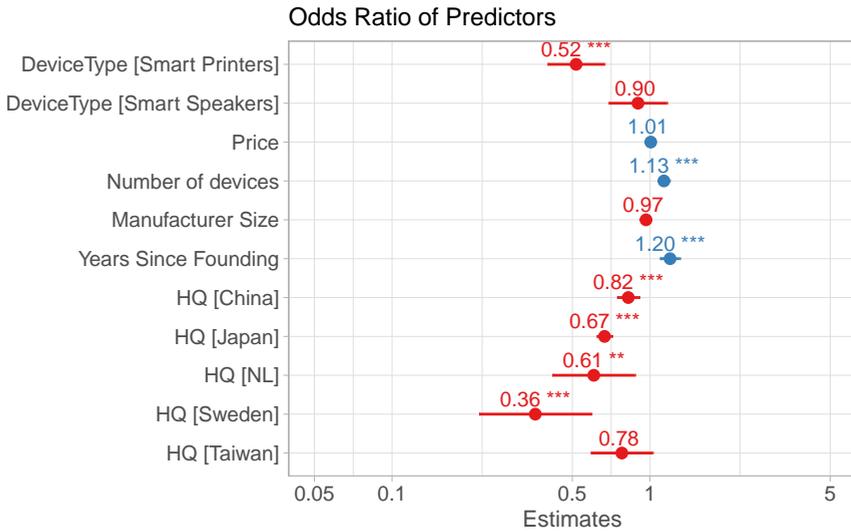
**Figure 4.7:** Proportion of feature changes per smart speaker manufacturer in each temporal trend

S&P, in 30 out of 60 S&P features, the proportion would be 0.5. Using the proportion of positive features allows us to compare the results across the three device types irrespective of the differences in the number of feature tests per device type. Additionally, compared to S&P rating, this allows for more transparency by eliminating the effect of the weights from the analysis. We use a Beta Regression Model which is best suited for proportions between 0 and 1.

We consider six predictor variables: two at the device level and four at the manufacturer level that have been identified in prior work as influencing an IoT devices' S&P [183]. At the device level, we include the categorical variable for device type and the device price. At the manufacturer level, we include the number of devices from the manufacturer in our dataset as a proxy for the diversity of device models they produce. We also account for the manufacturer's size (measured by the number of employees), the number of years since its founding, and the location of its headquarters. To standardise our analysis, we use scaled and centred values for variables such as price, number of devices, manufacturer size, and years since founding. Scaling ensures that all variables are on a similar scale, preventing any single variable from dominating the analysis. Centring adjusts the variables so their means are zero, which makes the intercept more meaningful and interpretation easier. Multi-collinearity tests showed that none of the variables were highly correlated, confirming their suitability for inclusion in the model.

We added the six variables in a step-wise forward manner, starting with the intercept only model and then adding the other variables, resulting in a total of seven models as shown in Table C.1 in Appendix C.1. We evaluated different orders for adding the variables in the model but found no differences based on the ordering. To assess model fit, we used the Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC),

and Log-likelihood, based on established guidelines from the literature [42, 65, 100, 120]. The AIC helps balance model accuracy and complexity, reducing the risk of overfitting, while the BIC places greater emphasis on simplicity by applying a stronger penalty for model complexity. Thus, lower AIC and BIC values indicate a better-fitting model. The Log-likelihood measures how well the model accounts for the observed data, with higher values suggesting a better fit. Based on these criteria, we identified Model 6 to be best fit among the models in Table C.1 (Appendix C.1) since it has the lowest AIC and BIC values and the highest log-likelihood among all the models.



**Figure 4.8:** Odds-Ratio of the Predictors from the Beta Regression Model. For categorical variables, the following reference categories were used: DeviceType = IP Cameras, HQ = USA.

Figure 4.8 shows the odds ratio of each of the predictors. The odds ratio denotes how a one-unit change in a predictor variable affects the odds of observing a higher or lower proportion of positive S&P results. For instance, the device type of smart printers has an odds ratio of 0.52. This indicates that, all else being constant, a smart printer has 48% lower odds of having positive S&P features compared to the reference category of IP cameras. The estimates for the other two device-level variables — the device type of smart speakers and device price — are not statistically significant.

At the manufacturer level, three variables are statistically significant: number of devices, years since founding and manufacturer headquarters; the size of the manufacturers has no significant effect. We observe that, all else being equal, a one unit increase in the number of devices or the years since founding is associated with a 13% and 20% increase, respectively, in the odds of having positive S&P features. This suggests that a manufacturer with greater experience, whether measured by age or by having released more models, tends to have higher odds of positive outcomes in the S&P tests. With regards to headquarters location, we find that a headquarter located in China, Japan, the Netherlands or Sweden decreases the odds of having positive results in S&P tests compared to the

reference category of having headquarters in the United States. Specifically, the odds are 12% lower for China, 33% lower for Japan, 39% lower for the Netherlands and 64% lower for Sweden. This indicates that IoT devices from manufacturers with headquarters in these countries are less likely to have devices with positive S&P results compared to devices from manufacturers based in the United States. This could be due to various factors such as differences in regulations, enforcement regimes, manufacturing practices, or organisation culture in different regions.

## 4.8. DISCUSSION

The main objective of our study was to determine whether manufacturers have improved the S&P of IoT devices over the years. Overall, our analysis finds limited evidence of improvement and uncovers a surprising trend. There is widespread inconsistent deployment of S&P features in successive IoT device models from the same manufacturer. Our analysis shows that even basic and critical S&P features like software update support are sometimes removed after being included only to be reintroduced later. We observe such fluctuating, inconsistent S&P feature deployments across all device types and manufacturers, in varying degrees.

These inconsistencies bring to focus the crucial and often overlooked role of the development process of manufacturers in determining S&P outcomes. A lack of structured development process with checkpoints at critical junctures to assess the S&P requirements could be one of the reasons for the inconsistencies. These manufacturers might face resistance towards implementing these checks due to entrenched organisational norms and cultures [219]. Even within a structured and established S&P development process, the underlying complexities of the IoT supply chain can impact S&P at the manufacturer level [15] like technical and contractual limitations in SDKs provided by chip vendors. Differences in the S&P priorities and feature implementations across different product development teams could also contribute to these inconsistencies. Overall, these indicate a lack of coherent S&P development framework at the manufacturer level.

The results of the Beta Regression Model also highlight the crucial role of the manufacturers. It shows that the proportion of positive S&P features in an IoT device is influenced by the location of the headquarters of the manufacturer and their experience both in terms of the number of models produced and years since founding. Contrary to other studies that suggest price of the device influences S&P [223], our model results show that device price has no significant influence when manufacturer level attributes are taken into account.

The importance of manufacturer experience emphasises the need for targeted support for manufacturers, especially smaller and newer ones who lack the needed experience. Such support can help them establish and maintain consistent S&P development practices and processes. Moreover, while principles like security by design - recommended in policy standards - are valuable, manufacturers also need practical guidance to operationalise these principles within their development process [110]. Additionally, the differences in S&P trends across various IoT device types highlight the necessity for tailored support specific to the needs of the development process of each device type, rather than a one-size-fits-all approach. Manufacturers in need of an overhaul to the S&P development processes can be suggested adequate mechanisms to overcome the organi-

sational resistance to change and implement mechanisms for continuous improvement. Similarly, manufacturers with an inconsistent S&P implementation across devices can be helped with creating a robust, structured approach to incorporating S&P within their development processes.

Overall, our results show that despite the introduction of various regulations and legislation, there has been little improvement in the S&P of the three IoT device types analysed in our study. Interestingly, our model results suggest that manufacturers headquartered in the US which operates under a less stringent regulatory environment compared to Europe, have higher odds of producing devices with more robust S&P features. This could be attributed to stricter enforcement actions [58] or the prevalence of legal repercussions, such as class action lawsuits, which are more common in the US [91]. These findings highlight the limitations of relying solely on market-driven regulations, which can lead to inconsistent S&P feature implementation based on the target market. Promoting awareness among manufacturers about codifying S&P best practices [198] into their development processes could help ensure consistency and reinforce existing regulations.

4

#### 4.8.1. LIMITATIONS

Due to our contractual agreement with the ECA, we cannot disclose the partner details, the test specifics, or names of the manufacturers and devices. While this limits the verifiability of the test results, we would like to highlight the strong reputation and ethical standards upheld by European consumer associations which lends a high level of confidence in the integrity and reliability of the testing process. Furthermore, our feature-level analysis specifies exactly which functionalities have remained stable or changed, offering valuable insights even in the absence of full test disclosure.

Our analysis is based on a limited dataset of IoT device S&P ratings from three device types and 23 manufacturers, which may not represent all manufacturers or device types. However, since consumer associations select popular devices for testing [223], we believe the analysis reflects trends in widely used IoT devices in Europe. It may not, however, capture trends in niche or emerging IoT device types. Another limitation is the accuracy of release dates, which were collected from various sources and may not always be precise. While ECA tests have remained consistent over time, there may be minor lab-specific variations that might introduce slight inconsistencies. Nevertheless, since we focus on broader trends, our findings remain robust to such variations.

Our analysis focuses on S&P features assessed by the ECA, which include a wide range of common S&P attributes in line with regulatory best practices. This implies that our findings provide valuable insights into most commonly recommended S&P features but may not capture trends in more specialized or emerging S&P aspects. Finally, we acknowledge that there are factors influencing IoT S&P features like supply chain dependencies, firmware practices, or third-party integrations that we were not able to include in our model. Despite this, the statistical significance of manufacturer-level factors in our analysis emphasizes the role of organisational processes in shaping S&P outcomes, extending beyond purely technical measures. Future research could explore these factors to provide a deeper understanding of the broader IoT S&P ecosystem.

## 4.9. CONCLUSION

In this paper, we analysed the evolution of S&P features over the past decade in 428 IoT devices from 23 manufacturers, focusing on three widely used and mature IoT device types: IP cameras, smart printers, and smart speakers. Our findings indicate that while IP cameras maintained consistently high S&P ratings, smart printers exhibited lower ratings with a slight declining trend, and smart speakers had the lowest ratings with no clear temporal pattern. At the manufacturer level, only a minority (3 out of 23) demonstrated improvement, while the majority (12 out of 23) maintained stable S&P ratings, and the remaining eight showed no clear trend in the ratings.

Our analysis also uncovers a surprising trend of inconsistent deployment of S&P features in subsequent device models of the same manufacturer. We find that the stability in ratings obscures such underlying inconsistencies. This highlights a need to help manufacturers operationalise S&P best practices within their development processes to ensure a more systematic and coherent focus on S&P across the development process.



# 5

## AVAILABILITY OF IOT UPDATE INFORMATION ON ONLINE STORES

*Security updates are essential for protecting IoT devices, yet consumers often lack reliable information about how long devices will be supported. We conduct the first large-scale study of update duration disclosures in the European market, analyzing 26,898 product pages across local retailers, EU Amazon sites, and Temu. Disclosure varies sharply: Dutch retailers, subject to regulatory oversight, list update durations for up to 92% of devices, while Amazon provides such information for fewer than 1% and Temu for none. For smart TVs, where EU rules mandate disclosure, coverage is higher but still inconsistent. Stated update durations vary between one and eight years, with smart TVs generally receiving the longest support. Comparing stated support durations across retailers, manufacturers, and the EU's central product database, we find widespread contradictions, with retailers often understating support relative to manufacturers. These inconsistencies limit the effectiveness of transparency mandates and risk misleading consumers. Our findings show that regulation can improve visibility, but only robust enforcement and standardized disclosure mechanisms ensure accurate and trustworthy information.*

### 5.1. INTRODUCTION

Many Internet-of-Things (IoT) devices are released with weak security postures, making them attractive targets for cyberattacks. Even when manufacturers equip their IoT devices with robust security features, they remain susceptible to vulnerabilities discovered after release [16, 104]. Without timely security updates, these vulnerabilities can be exploited at scale, turning millions of interconnected devices into entry points for attackers.

Security updates ('patches') are therefore not just a best practice, they play a critical role, providing the means to retroactively patch security vulnerabilities in IoT devices. Recognising this, the provisioning of security updates has become the focus of legislation and policy efforts in different jurisdictions requiring IoT manufacturers to provide more transparency. Consumers should be informed at the time of purchase about how long

a device will receive security updates. Update duration is a critical focus because, like a product warranty, it can be clearly disclosed at purchase and factored into consumer decisions, empowering buyers and incentivizing manufacturers to extend support.

The US Cyber Trust Mark [79], a voluntary IoT device labelling initiative, expected to take effect in late 2025, includes update support duration as part of its label. Similarly, in the UK, the Product Security and Telecommunications Infrastructure (PSTI) Act places a legal requirement on IoT device manufacturers to provide information to consumers about how long their products are supported by security updates [66]. The Cyber Resilience Act (CRA) in the European Union mandates manufacturers to disclose the update support duration for their products. It also requires that security updates be provided for a device's entire expected use period, and at least for five years [74]. Another EU Regulation, on Energy Labelling [3], provides an obligation on manufacturers to provide update duration information for smart TVs, in addition to the energy ratings. Manufacturers have to submit this information to a central, consumer accessible, database called the European Product Registry for Energy Labelling (EPREL).

Online stores are a critical interface where consumers engage with product information. If regulatory and voluntary efforts to inform consumers are having an effect, update duration should become visible on product pages. Large stores also have more leverage than consumers to get clear disclosures from manufacturers about their update support for a product. A complementary regulatory strategy is to require online stores to include this information in their product pages. In the Netherlands, this strategy has been in place since 2020. In that year, the Netherlands Authority for Consumers and Markets (ACM) [26] required major online retailers to provide update support duration information.

Despite these regulations emphasising provisioning of update support duration, there is very little empirical data on the presence of this information in the marketplace or on, where this information exists, what update durations manufacturers are promising. The US Federal Trade Commission (FTC) recently conducted a small-scale study on 184 devices and found that nearly 89% of manufacturer product webpages failed to disclose software update durations [182]. The only academic study we are aware of reported that update duration was available for 20% of the 417 IoT devices that were analyzed [224]. In sum, there is a lack of large scale systematic evaluation on the availability of update support durations and on the length of the durations currently promised. This gap makes it difficult to assess market transparency and hinders regulators and researchers from evaluating the effectiveness of current efforts to improve security update support.

In this paper, we address this gap through three research questions, focusing on online stores where update information can be collected at scale, is most visible to consumers, and most relevant. Since purchases typically form a contract with the store rather than the manufacturer, what the store discloses matters. Our first question is: *To what extent do online stores in the EU provide information on the update support duration of IoT devices?* To answer, we manually reviewed 4,600 product pages for five device types across 58 EU online stores: regional retailers, country-specific Amazon sites, and Temu.com. The data was collected between Dec 2024 and May 2025. Of the 22 regional retailers, 16 provide no update information, 5 disclose it only for smart TVs, and Dutch retailers list durations for all device types. Among global players, Temu.com shows none, while Amazon's EU sites list durations for only some devices.

Our second research question was *What percentage of devices include update support duration, and what is the distribution of the durations?* Building on the results from the manual analysis, we scraped 22,298 product pages: all device types on EU Amazon sites and two Dutch sites, plus smart TVs from regional retailers (the only cases with update data). Amazon rarely provides this information (0.4%), despite having the update support duration field on every page. Dutch retailers disclose far more (15–92%), and smart TVs in particular show higher coverage (20–98%). Stated durations range from one to eight years, with smart TVs typically receiving the longest support.

Finally, we compare update information across sources—retailer sites, manufacturer websites, and for smart TVs, the EPREL database—asking: *How consistent are disclosed update durations across sources?* We focus on regional stores for smart TVs and on Dutch stores for all device types, as Amazon provides too little data. Moreover, Dutch retailers and smart TVs are a special case due to regulatory intervention and an in-depth analysis can help identify ways to improve the status quo elsewhere. We supplemented our dataset with manufacturer pages for 250 devices and EPREL entries for smart TVs. The results show major inconsistencies between retailers, manufacturers, and the central database, casting doubt on the overall reliability of update support information. In sum, our contributions are as follows:

- We present the first large-scale analysis of update information provisioning, focused on European online stores. We manually review 4,480 product pages for disclosures and scrape an additional 22,298 pages for characterizing the statements of update durations.
- We find empirical evidence for the impact of two regulatory initiatives, for Dutch online stores and for smart TVs across several countries, but also characterize significant gaps in compliance.
- For a smaller sample, we provide a comparison of update duration availability across three sources. Surprisingly, online stores have higher availability of update information than (between 15.2% and 91.5%) manufacturer websites (between 18% and 50%), while the EPREL database has update information for around 80% of the TVs analysed.
- For the same device models, we identify conflicting update information between stores, manufacturers and EPREL database, highlighting inconsistencies in the disclosures to consumers. We offer recommendations for different stakeholders to better support consumers making more security-conscious IoT purchase decisions.

## 5.2. RELATED WORK

We discuss two main areas of related work. First, earlier studies that analyze whether information is disclosed about how long an IoT device will receive updates. And second, research on the role of updates in consumer decisions.

There is surprisingly little work on the availability of update support information. The FTC checked the manufacturer webpages for 184 devices and found that 89% of these pages did not disclose the update support duration. An academic study on the

relationship between market sales and security features of IoT devices checked the availability of update support information for a sample of 417 devices across two online stores [224]. Similar to the FTC, it reported that valid information is present in only a fraction of the product pages. A study by Privacy International [106] checked the websites of 21 manufacturers across five categories of devices (smartphones, personal computers, gaming consoles, tablets and smart TVs). Some manufacturers provide clear information on update support duration, but for most the information is either unclear or missing altogether.

These studies are based on relatively small samples of a few hundred products at most and only evaluate manufacturer disclosure. In contrast, we present the first large-scale analysis of disclosure and a characterization of the update support durations stated for IoT devices. In all, we analysed 26,898 product pages across 55 online stores in Europe. While we also look at manufacturer disclosures, we put more emphasis on online stores, since that is where most consumers will interface with this information. Also, a purchase means the consumer is entering into a legal contract with the store, not with the manufacturer directly.

5

The second area of related work has focused on understanding consumer preferences around security updates for IoT devices. Some studies [133, 222] show that users expect that security updates should be offered for a reasonable period and that this reasonable period would vary based on the device category. A consumer study [105] shows that 31% of users expect that internet connected home devices will receive regular security and system updates for between 2 and 10 years. In a survey of 412 smart home users [95], the researchers find that while a majority of the users believed that updates were important for smart devices, they were less concerned about when the manufacturer would stop supporting their device. In contrast, a study into the ‘Willingness to Pay’ [163] found that consumers are willing to pay more for lifetime software updates than for 5 years of updates. In a more qualitative study [94] with interviews of 40 smart home users about their experience with updates. They note that users rarely associate updates with security, are unclear on the update processes, and are concerned about manufacturers discontinuing support due to the dynamic nature of the smart home market. Another study [157] analysed the influence of availability of security updates into consumer purchase decisions by conducting a survey with around 1400 participants. It finds that availability of update support duration information can explain between 8% and 35% of the variance in consumers purchase choices.

Our study does not analyze consumer preferences or purchasing behavior, but we do build on this prior work in the sense that we study the provisioning of the information at the time of purchase, knowing that it influences consumer choice. It provides the motivation for our study to focus on online stores, which is where consumers engage with product descriptions, rather than manufacturer websites. (Though we also study the latter for a specific sample.)

### 5.3. METHODOLOGY

To address our first research question on the level of provisioning currently present in EU online markets, we manually analysed the online stores and identified the stores which provide update information. For our second research question on characterising the

update support duration currently stated by retailers for different IoT device types, we do large scale data collection and analysis using web scraping and extraction techniques. In this section we detail the methodology followed for both.

For both research questions, we considered five IoT device types that have been shown to be popular in the consumer IoT landscape – IP cameras, smart printers, smart speakers [132, 224]. We scoped our analysis to the 30 countries in the European Economic Area (EEA) at the time of writing this paper: Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Netherlands, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain and Sweden.

### 5.3.1. WEB STORE DATA COLLECTION

To evaluate the update provisioning comprehensively, we wanted to systematically identify and analyse online stores popular in each of the countries. To do so, we used Similarweb, a digital intelligence platform that provides market research and trends for websites and apps. Based on the results from Similarweb, we identified three types of online stores: local or region-specific platforms tailored to the domestic market, the Amazon store popular in that region and Temu, the Chinese e-commerce platform that is popular across Europe. In all EU countries, Temu was ranked in the top 5.

**Table 5.1:** Research Aims and Dataset Overview

Research Question	Research Aim	Websites and Device types	Count
RQ1	Evaluating Update provisioning on webstores	All five device types on: - 22 regional websites - 7 Amazon websites - 29 regional versions of Temu	4,480 product pages
RQ2	Quantifying update information presence and duration	All five device types on: - 6 Amazon websites (.de, .fr, .it, .nl, .es, co.uk) - 2 Dutch Websites (bol.com, coolblue.nl) Only smart TVs on: - 3 regional websites (kaup24.ee, 1a.lv, pigu.lt)	22,298 product pages
RQ3	Comparing across sources	Manufacturer websites + EPREL Database	272 devices

Out of the 30 countries considered, Similarweb data was unavailable for 9. In these cases, we supplemented our analysis with Google search to identify the most prominent

local online store. We searched for '*popular e-commerce websites in <countryname>*' and, based on the results, manually identified the leading local online store based on consumer or e-commerce blogs. We were not able to identify a local online store in eight of the 30 countries since the websites mentioned in the google results were not local to the region.

At the end of this step, we had a list of 22 local online stores and 7 Amazon websites that were popular across the EEA countries. Together with Temu, these 30 websites provide a starting point for our analysis. While the URL for Temu remains the same across Europe, the website allows users to select from 29 of the 30 EEA countries as their region and based on the region selected, the offerings vary. Although we analysed each of these 29 regional versions individually, we count Temu as a single website — rather than 29 separate ones — for simplicity.

### PRODUCT CATEGORY PAGE IDENTIFICATION

In the next step, we identified the relevant product category page - the page that contains all listings for a particular product - for each of the 30 websites. We identified the product category pages for all five IoT device types: IP cameras, smart printers, smart speakers, smart TVs and smart watches. This approach is better than searching for categories like 'IP Cameras' since search results are often cluttered with auxiliary products (e.g., IP camera wall mounts, batteries, dummy cameras, etc.). When no dedicated category existed, we used the closest match and applied filters — for example, selecting smart speakers from the general speaker category by filtering for WiFi connectivity.

5

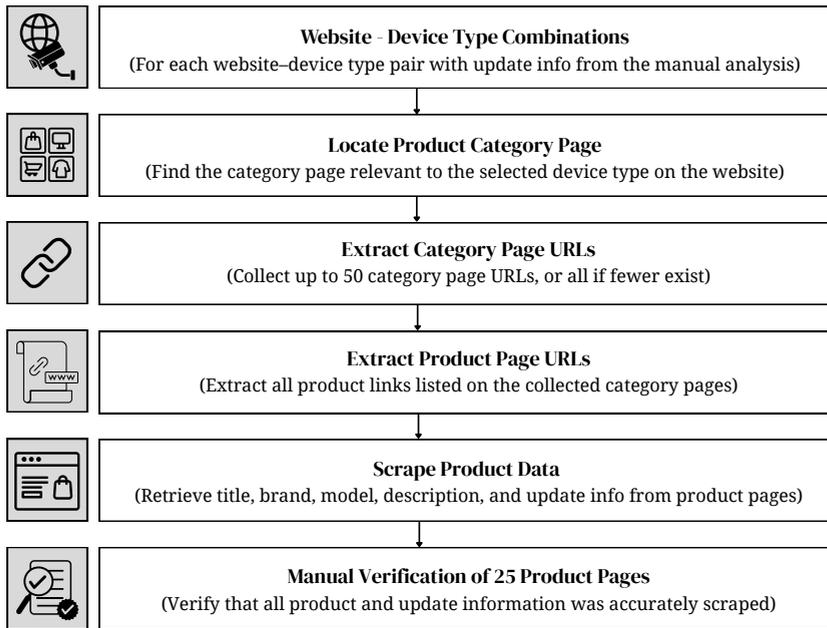
### 5.3.2. EVALUATING PROVISIONING: MANUAL ANALYSIS OF ONLINE STORES

To answer our first research question, we manually evaluated the first twenty product links for each of the 5 device types, on each of the 30 websites to identify whether any update duration related information is available. We use product pages as the unit of analysis because they are the primary point of compliance for disclosing the update information. Although some devices have multiple pages due to different colours, sellers, bundles etc., we treat each page individually to capture all instances of disclosed information and avoid missing variations in update support details. The results of this evaluation feed into the next step.

### 5.3.3. CHARACTERIZING UPDATE DURATIONS: LARGE SCALE ANALYSIS

To answer our second research question, we scraped the websites and device type combinations that we identified in the previous step as containing update support duration information. We scraped the product links from the first fifty pages of each of the product category listings. Each page typically contained around 20-30 product links. Next, we opened each of these product links, saved the HTML pages and scraped the product data. We used BeautifulSoup, a Python library for parsing HTML, for scraping and extracting the product data – title, price, brand, and model details (where available), along with the product information or any specific fields identified in the previous step as containing update support duration information.

Next, we manually verified if the scraped information was accurate for 25 random listings of each device type and website combination, analysing 1,075 product pages overall. This step was crucial to make sure that the scraping scripts did not miss any



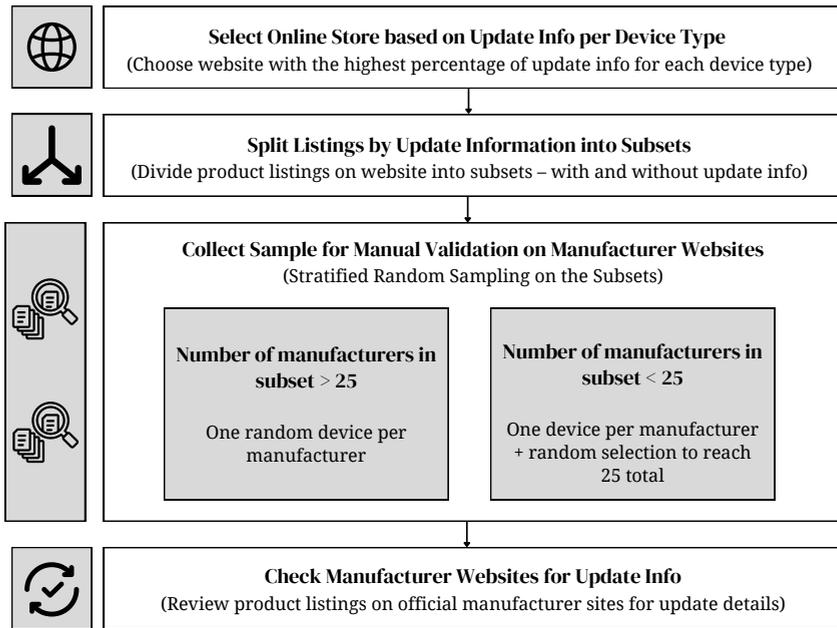
**Figure 5.1:** Workflow for Large Scale Analysis

update information.

#### 5.3.4. COMPARISON ACROSS SOURCES

To answer our third research question, we compared the update information between two online retail stores, between retailers and manufacturers and for smart TVs, also with the centralised database. We examined manufacturer websites since regulations emphasize their responsibility to provide update support details. Our objective with this analysis was twofold. First, we wanted to compare whether the durations stated by the manufacturer were consistent with those stated on the online stores. Second, we aimed to check the overlap between devices with update information on online stores versus the manufacturer’s website. Our hypothesis was that there would be a high degree of overlap, since if the information is available in one source, it should be straightforward to populate it in the other as well.

We used a stratified sampling strategy to select devices for manual validation. For each device type, we chose the online store with the highest percentage of products listing update durations. Listings were split into two groups (with and without update information), and at least 25 devices were randomly sampled from each, ensuring at least one per manufacturer. If fewer than 25 manufacturers existed, additional devices were randomly added; if more than 25, one device per manufacturer was randomly selected and 25 devices were randomly sampled from this pool. This yielded at least 50 devices per type with diverse manufacturer representation.



**Figure 5.2:** Workflow for Manufacturer Website Analysis

To find manufacturer-published update information, we searched Google using the manufacturer and model to find the product page on the manufacturer's website. If the page lacked update details, we repeated the search with the term 'update support duration' to capture cases where the information was presented in a separate page. This part of the analysis was conducted in March 2025.

For the centralised database, we compared the smart TV details on the website<sup>1</sup> and got the details where available.

## 5.4. PROVISIONING OF UPDATE DURATION INFORMATION ON EU ONLINE STORES

In this section we present the results of the first research question on the degree of update provisioning on EU online stores, starting with the local stores, then the Amazon stores and finally Temu.

### 5.4.1. LOCAL ONLINE STORES

Of the 30 EEA countries, 22 had a local online store ranked in the top five on Similarweb based on popularity; the remaining eight did not. In two cases, a single store served two countries (bol.com for Belgium and the Netherlands, skroutz.gr for Cyprus and Greece). As listings can vary by country, we set the relevant country or delivery location and

<sup>1</sup><https://eprel.ec.europa.eu/screen/product/electronicdisplays>

analysed them independently, resulting in 20 unique local online stores across the 22 countries.

**Table 5.2:** Results of the Manual Analysis of Local Websites per Country

Country	Website	IP Camera	Smart Printer	Smart Speaker	Smart TV	Smart Watch
Belgium	bol.com	●	●	●	●	●
Bulgaria	emag.bg	○	–	○	○	○
Croatia	emmezeta.hr	–	○	–	●	○
Cyprus	skroutz.gr	○	○	○	○	○
Estonia	kaup24.ee	○	○	○	●	○
Finland	verkkokauppa.com	○	○	○	○	○
France	cdiscount.com	○	○	○	○	○
Germany	otto.de	○	○	○	○	○
Greece	skroutz.gr	○	○	○	○	○
Hungary	emag.hu	○	○	○	○	○
Iceland	elko.is	–	○	○	○	–
Latvia	1a.lv	○	○	○	●	○
Liechtenstein	ricardo.ch	○	○	○	○	○
Lithuania	pigu.lt	○	○	○	●	○
Netherlands	bol.com	●	●	●	●	●
Norway	clasohlson.com	○	○	○	–	○
Portugal	worten.pt	○	○	○	○	○
Romania	emag.ro	○	○	○	○	○
Slovakia	mall.sk	○	–	○	○	○
Slovenia	mimovrste.com	○	○	○	○	○
Spain	elcorteingles.es	○	○	○	○	○
Sweden	clasohlson.com	○	○	○	–	○

**Legend:** ● Update info present ○ No update info – Not sold

Table 5.2 lists the local websites analysed for each country and the availability of update information on the site. Of the 20 websites, five did not sell one or more of the device types considered, indicated with (–). Cases where a device type was sold but lacked update information are marked (○), and cases with update information are marked (●). Most websites and device types had no update information; only 7 of the 22 countries had local sites showing update support for at least one device type.

The table also highlights two patterns possibly driven by policy interventions. First, the Dutch site bol.com (also serving Belgium) provides update information for all device types, likely due to a requirement by the Netherlands Authority for Consumers and Markets (ACM) [26]. Second, update information for smart TVs appears on about half of local stores, possibly driven by an EU energy labelling directive that mandates reporting software and firmware support periods. We delve deeper into these in the next section on characterizing the update durations.

### 5.4.2. AMAZON STORES

Next, we analyzed Amazon stores, focusing on the local storefronts for EEA countries. Of the 30 countries, 14 did not have Amazon among the top five popular sites. For the remaining 16, some shared the same regional Amazon website—for example, amazon.es

was popular in both Spain and Portugal. We therefore analyzed seven Amazon sites: amazon.co.uk (Ireland), amazon.com (Bulgaria, Croatia, Hungary, Norway, Romania, Sweden), amazon.de (Austria, Denmark, Germany), amazon.es (Portugal, Spain), amazon.fr (Belgium, France), amazon.it (Italy), and amazon.nl (Netherlands). Delivery addresses were set locally to obtain relevant results.

Our manual analysis of 20 product links per device type showed that all the European Amazon websites have update support duration provisioning. In contrast, although amazon.com was popular in six EEA countries, it had no update information for any of the 100 devices analysed. We characterize the percentage of devices on the European Amazon websites with the update information and the duration specified in the next section.

#### 5.4.3. CHINESE ONLINE STORE - TEMU

Next, we analysed Chinese Temu websites that are popular across Europe. As mentioned earlier, although Temu does not have region-specific domains, it provides an option to change the country or region on its website. Of the 30 EEA countries, only one — Liechtenstein — was not available in the list of regions. We analysed the remaining 29 regional versions of Temu for three types of IoT devices: IP cameras, smart speakers, and smart watches. At the time of our data collection, smart TVs and smart printers were not sold on Temu. We conducted a manual analysis of the first twenty product listings for each of the three device types across all 29 websites and found no update information on any of them.

### 5.5. QUANTIFYING UPDATE SUPPORT DURATIONS

In this section, we answer our second research question. We quantify the presence and distribution of update support duration information across the websites and device types identified in the previous section as containing update information. The websites that have the update information for device types other than smart TV are bol.com and the six Amazon websites. To enable more comprehensive analysis, we also included another popular Dutch online store coolblue.nl [179], which would also help us better evaluate the effect of the Dutch intervention [26].

On these eight websites, we scraped product links for IP cameras, smart printers, smart speakers and smart watches from the first fifty pages of product listings (or all pages if the total was less than 50). We categorise the presence of update information into three categories, the update field was (a) absent, we classify this as field unavailable, (b) present but says Unknown or Not applicable or (c) specifies a duration (e.g., At least 60 months after the introduction date) or date (e.g., Guaranteed software updates until 15 Sep 2031). bol.com and coolblue.nl list the duration while the Amazon websites list a date.

Table 5.3 reports the fraction of pages in each category across all the seven websites. On bol.com, the percentage of product links with update information ranges from 33.1% to 64%. The other Dutch website, coolblue.nl, has update information for 17.8% to 91.5% of product links, with no smart printers having update support information. The distribution of update duration across both the bol.com and coolblue.nl ranges from 1 year to 7 years across all the device types except smart TV. We discuss smart TV separately

**Table 5.3:** Percentage of product links in each update status category across EU websites for all device types except smart TVs

Website	Update Info Field	IP Camera	Smart Printer	Smart Speaker	Smart Watch
bol.com	<i>Total Count</i>	1190	882	127	1152
	% Unavailable	0.0	5.4	0.0	0.0
	% Says Not Applicable	38.3	61.5	44.8	36.0
	% Specifies Duration	61.7	33.1	55.2	64.0
coolblue.nl	<i>Total Count</i>	248	111	79	319
	% Unavailable	14.9	100	58.1	6.3
	% Says Unknown	8.1	0.0	24.1	2.2
	% Specifies Duration	77.0	0.0	17.8	91.5
amazon.de	<i>Total Count</i>	738	882	26	954
	% Unavailable	73.6	7.7	23.1	15.2
	% Says Unknown	26.0	92.3	76.9	84.4
	% Specifies Date	0.4	0.0	0.0	0.4
amazon.es	<i>Total Count</i>	896	470	33	979
	% Unavailable	81.6	98.1	21.2	27.0
	% Says Unknown	18.4	1.5	75.8	73.0
	% Specifies Date	0.0	0.4	3.0	0.0
amazon.fr	<i>Total Count</i>	919	552	32	454
	% Unavailable	73.8	98.0	18.8	99.6
	% Says Unknown	25.8	1.8	81.2	0.4
	% Specifies Date	0.4	0.2	0.0	0.0
amazon.it	<i>Total Count</i>	797	818	32	410
	% Unavailable	84.4	97.4	28.1	95.9
	% Says Unknown	15.3	2.6	71.9	4.1
	% Specifies Date	0.3	0.0	0.0	0.0
amazon.nl	<i>Total Count</i>	1071	371	33	405
	% Unavailable	65.6	1.1	0.0	19.3
	% Says Unknown	33.1	94.9	100.0	79.8
	% Specifies Date	1.3	4.0	0.0	1.0
amazon.co.uk	<i>Total Count</i>	1236	721	1167	1203
	% Unavailable	73.5	9.3	69.6	20.8
	% Says Unknown	26.1	90.7	30.4	79.2
	% Specifies Date	0.4	0.0	0.0	0.0

in the next section.

In contrast to the Dutch websites, the vast majority of devices on the Amazon websites lack specific update information: 70.4% of product links have no update support duration field, while 29.3% display the field but list it as Unknown. Only 0.3% of all links specify a date. Across the four device types and six EU Amazon websites, only 51 product links specify a date. We started our data collection in 2024, but 18 product links (35.3%) listed a date in 2023 or earlier in the 'Guaranteed software updates until' field. Two links state a date in 2083, while the rest range from 2026 to 2031, with 13 April 2030 being the most common (41.2% or 21 product links).

Device type differences exist: IP cameras and smart watches have higher percentage of availability across both the Dutch websites and the EU Amazon websites. Within the Amazon websites, regional differences are also notable. The Dutch site, amazon.nl, performs best, with 4% of smart printers and 1.8% of smart TVs specifying dates, reflecting broader local compliance, though overall adherence to ACM requirements remains low (0–4%). In contrast, amazon.fr and amazon.it consistently show 0–0.4% of products with update information. Despite the UK PSTI Act [66], amazon.co.uk lists dates for only 0.1% of smart TVs and 0.4% of IP cameras, with no dates for other device types, suggesting UK consumers rarely receive useful update duration information despite policy and a dedicated field.

5

**Smart TVs** As noted in the Introduction, smart TVs represent a unique case under the Commission Delegated Regulation (EU) 2019/2013 [3], which introduced mandatory disclosure of software update support durations for electronic displays as part of the EU's revised energy labelling framework (effective March 2021). Smart TVs are the only IoT devices explicitly covered by this regulation, which requires product information sheets to state the minimum guaranteed availability of software and firmware updates alongside energy performance data. This requirement, linked to the ecodesign regulation for electronic displays (EU 2019/2021), supports EU policy goals on sustainability, transparency, and digital rights. It is intended to help consumers make informed decisions about device longevity, particularly software support for connected products [3, 4].

Our analysis found that five of the 22 local online stores and all EU Amazon websites included update information for smart TVs, though in varying formats. Retailers in Croatia (emmezeta.hr), Estonia (kaup24.ee), Latvia (1a.lv), Lithuania (pigu.lt), and the Netherlands (coolblue.nl) provided it in PDF product information sheets linked from their sites. On all six European Amazon websites, the sheets were available only as images, making the content unsearchable and less accessible. The Dutch retailer bol.com displayed the QR code with the energy label in a product carousel image, which directed consumers to the product's EPREL page.

Table 5.4 shows the percentage of links per website that contained update information. None of the websites provided the sheets for all listed products. Data from the Croatian site emmezeta.hr could not be collected due to instability: the site frequently failed to load and often redirected product links to "not found" pages.

We observe substantial variation in the percentage of devices with update support information across countries and websites. On local websites, less than 20% of smart TVs on Latvia's 1a.lv include the product information sheet compared to 81.5% of TVs in

**Table 5.4:** Percentage of Smart TVs with product information sheet

Website	Country	Percentage of product pages with product information sheet
kaup24.ee	Estonia	33.3%
la.lv	Latvia	19.8%
pigu.lt	Lithuania	81.5%
bol.com	Netherlands	97.6%
coolblue.nl	Netherlands	54.2%
amazon.de	Germany	91.8%
amazon.fr	France	46.4%
amazon.it	Italy	43.2%
amazon.nl	Netherlands	37.9%
amazon.es	Spain	50.4%
amazon.co.uk	UK	41.2%

Lithuania's pigu.lt. The Dutch websites also show a sharp difference. A little more than half (54.2%) of TVs on coolblue.nl have the update information compared to 97.6% on bol.com. Among the Amazon websites, Germany has the highest percentage of update information (91.8%) while UK has the lowest (41.2%).

Consistent with the regulatory mandate requiring updates for eight years, the most frequently stated duration in product information sheets was also eight years. On some websites, up to 93.5% of TVs specified an eight-year update period.

## 5.6. COMPARISON ACROSS SOURCES

In this section, we move beyond quantifying the presence of update information to analyse update durations by device type. We compare the two Dutch retailers, manufacturer websites, and the centralised smart TV database to address our third research question. Amazon is excluded from this analysis, as it provides update information for only a negligible percentage of products.

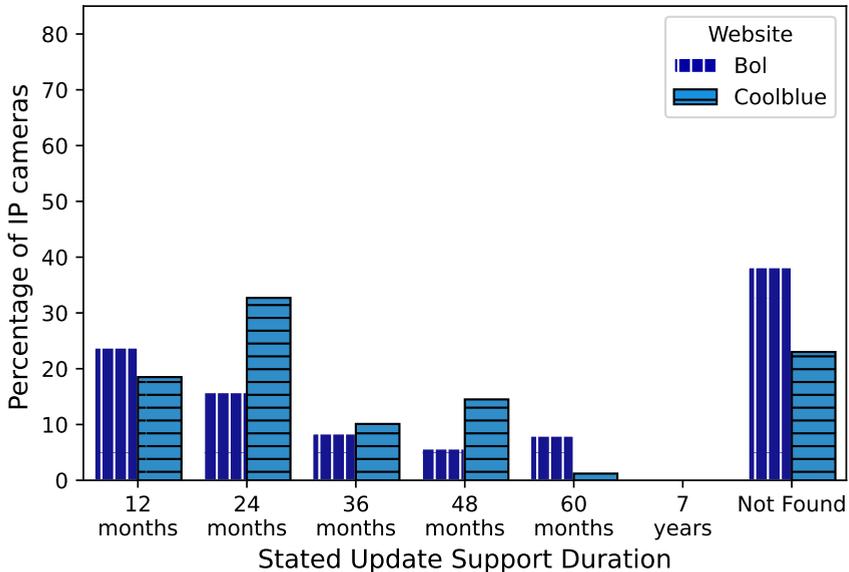
### 5.6.1. COMPARISON BETWEEN DUTCH RETAILERS

Between the Dutch retailers, we first compare the distribution of update support durations across all the product links. Next, we also compare the update durations for the specific device models that are available on both the websites. To do so, we first filtered brands available on both bol.com and Coolblue, then extracted model names from product titles using regular expressions (e.g., E340 for a Eufy camera, C200 for a TP-Link camera, or Vivofit Junior 3 for a Garmin smartwatch). Model names were manually verified and edited where needed. Duplicate listings which are common due to color, material, or bundling variations, were dropped. From the resulting unique sets, we selected models appearing on both sites for direct comparison. Because only a few devices included all three fields, we focused only on the stated update support period, and excluded the release year and month. While this approach may have overlooked some cases where the same model was sold on both but without a distinctive model name, it also ensures that the devices we did compare were indeed the same models. Since coolblue.nl did not have any update information, we only compare for IP cameras, smart speakers and smart watches. We analyse the smart TVs separately.

On both bol.com and coolblue.nl, the update support duration information is presented in terms of number of months from the date of introduction of the device under the fields 'Support with updates' and 'Guaranteed support with updates' on bol.com and coolblue.nl respectively. There are also two additional fields which state the month and year of introduction ('Introduction year' and 'Introduction month' on both websites). We discuss the results of our analysis for each device type and present the results for smart TV separately (in [subsection 5.6.3](#)) since that is a special case.

### IP CAMERAS

On bol.com, we collected data on 1,190 IP cameras from 247 brands. Of these, 61.7% (735 cameras) included update support information, but only 36.2% had all three key fields completed. Another 25.5% listed update support but lacked either the introduction year or month. In contrast, coolblue.nl offered had fewer listings, 48 cameras from 10 brands, yet about 80% included all three fields, showing more consistent update data despite the smaller range. As shown in [Figure 5.3](#), both platforms most often listed update durations of 12 and 24 months.



**Figure 5.3:** Comparison of update durations for IP cameras on bol.com and coolblue.nl

On Bol, 50 brands had at least three cameras, enabling intra-brand comparison. For 16 brands, the update duration was identical across models (standard deviation = 0), while the other 34 brands showed inconsistencies, with deviations up to 27.7 months. Price offered no explanation: expensive cameras were no more likely to receive longer support.

On Coolblue, only 7 brands had three or more devices, but variation was still evident. Ring (68 cameras) and Google (9) showed consistent support of 24 and 36 months,

respectively, while Eufy and TP-Link varied. Interestingly, we found a small but statistically significant negative correlation ( $r = -0.27, p < 0.01$ ) between price and update duration—contrary to the intuition that higher-priced cameras should be supported longer.

For the device model level analysis, we identified 38 models common to bol.com and coolblue.nl. Of these, 24 had update information listed on both websites. Half of these (12 models) showed identical information, while the other half (12 models) displayed conflicting durations, indicating significant inconsistency between the platforms. The discrepancies did not follow a clear manufacturer or brand pattern. All seven Ring cameras had inconsistent information, four of the Imou cameras were consistent while one was not, and among seven TP-Link cameras, three had information only on one website, three were consistent across both, and one was inconsistent.

SMART SPEAKERS

We analysed smart speaker listings on bol.com (127 devices) and coolblue.nl (79 devices) to assess the availability of update support information. On Bol, 55.1% of speakers included some duration data, with 34.6% having all three fields (duration, introduction year, month) completed. In contrast, coolblue.nl showed far lower availability: only 15.2% had complete information, while 58.2% had the update field missing and 24.1% marked it as Unknown.

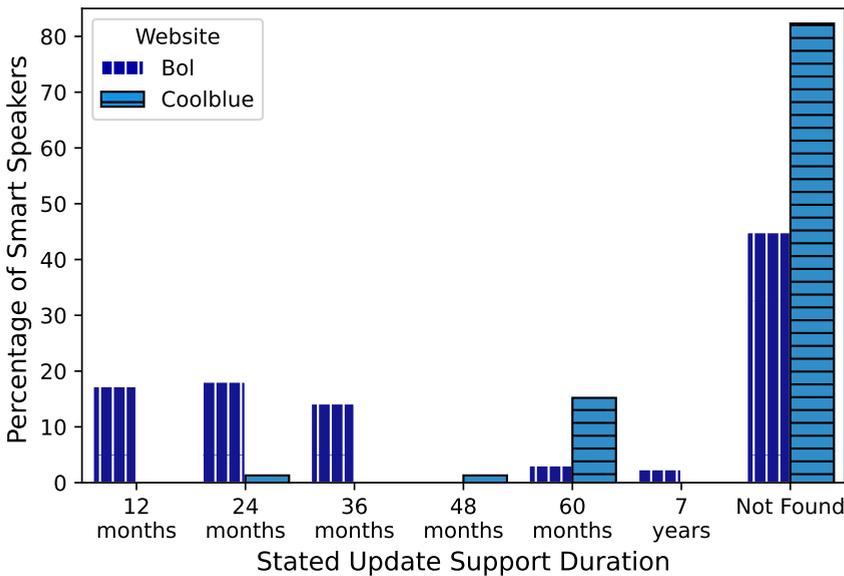


Figure 5.4: Comparison of update durations for Smart Speakers on bol.com and coolblue.nl

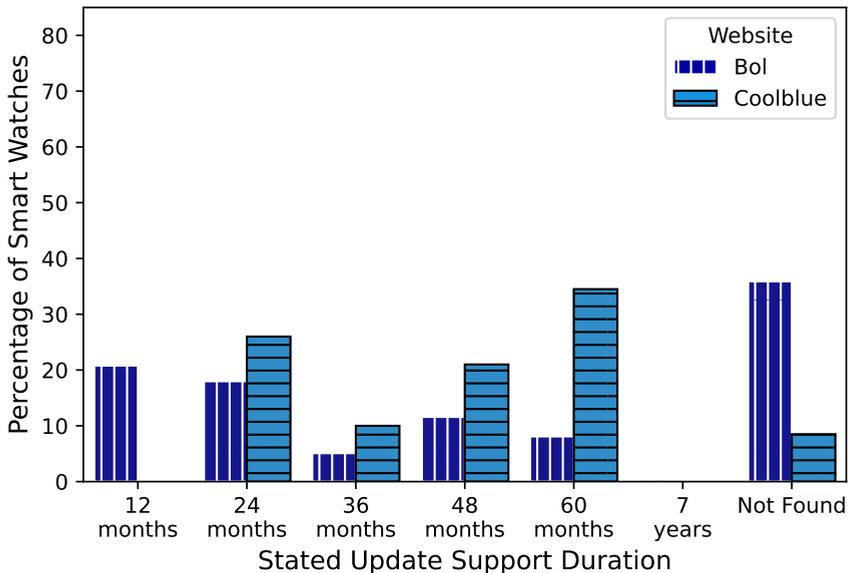
Figure 5.4 compares stated durations. On Bol, the most common values were 24 months (18.1%), 12 months (17.3%), and 36 months (14.2%), with a small fraction (2.4%) listing 7 years. Only two brands provided data for more than two devices: Kodak (8

speakers, all 12 months) and Bluesound (3 speakers, 36–60 months). On Coolblue, eight brands had more than two speakers listed, but only Denon provided update data. None of the others, including Apple, listed update support.

At the device model level, only ten models appeared on both websites. Of these, just two included update information on both, and in both cases the information conflicted: each speaker listed 24 months of support on bol.com but 60 months on coolblue.nl.

### SMART WATCHES

We also analysed 1,152 smart watches on bol.com and 319 on Coolblue. On Bol, 64% of watches had update support information, but only 10 devices (all from TCL) included all three fields. By contrast, coolblue.nl listed durations for 91.5% of watches, with nearly all (except two) also including the introduction year—but none the month. As shown in [Figure 5.5](#), both platforms reported durations ranging widely, with 24, 48, and 60 months the most common.



**Figure 5.5:** Comparison of update durations for Smart Watches on bol.com and coolblue.nl

Across 54 brands with more than two watches (covering 58.7% of devices), market concentration was lower than for printers. Within-brand variability was high: Apple ranged from 12–60 months (average 49.4), while Garmin averaged 41.8. A Kruskal–Wallis test confirmed significant differences across brands ( $H(676) = 476.13$ ,  $p < 0.001$ ). On Bol, we found a modest positive correlation between price and update duration ( $r = 0.356$ ,  $p < 0.01$ ), suggesting higher-priced watches tend to list longer support. Coolblue, by contrast, showed higher concentration: its 319 watches came from just 10 brands, with most listing a single consistent duration. Even Apple, Garmin, and Samsung showed consistency for the majority of models. Due to the smaller dataset, no correlation analysis was possible for Coolblue.

For the same device models of smart watches, we found 44 models common to both sites. Of these, 16 listed consistent information across platforms, while 18 listed conflicting durations. In addition to these cross-platform differences, we also observed contradictions within the same store. On bol.com, for instance, 30 listings of the Apple Watch Ultra were divided between 13 showing a duration of 24 months and 17 showing 48 months. Similarly, among 54 Apple Watch Series 10 listings, 53 stated 60 months while one listed only 24 months. Such inconsistencies are particularly noteworthy because Apple is a major brand with a reputation for tightly controlled product information. For consumers, this creates uncertainty about how long their devices will actually be supported, and raises concerns about the reliability of the update data published by retailers.

Taken together, these findings highlight a prevalent lack of consistency in update support information on bol.com and coolblue.nl, both between platforms and within the same retailer. In many cases, one store provided update information while the other did not, and when both did, the information often contradicted. This challenges the assumption that retailers simply act as passive channels for manufacturer-provided data. To better understand where these discrepancies originate, we contacted customer service on both platforms, posing as potential buyers. bol.com did not respond, while coolblue.nl stated that update durations are based on ‘internal case studies’ but could not share any publicly available sources. Since update durations vary even within the same brand, these internal case studies appear to apply at the model level rather than at the manufacturer level. In any case, the discrepancies raise important questions about the true source of these durations and about which durations, if any, are aligned with the claims of the manufacturers.

### 5.6.2. BETWEEN RETAILERS AND MANUFACTURER WEBSITES

The consistencies identified raises the question of how these durations compare to what the manufacturers state. To address this, we analyse the update information stated by the manufacturers. As discussed in subsection 5.3.4, we checked for manufacturer disclosure for two sets of devices per device type: one set of devices that do have update information on the Dutch stores and one set that does not. To check whether update information was available on manufacturer websites, we first did a google search to locate the product page for each device on the manufacturers website (e.g., Apple Watch Series 10). Next, we searched for pages containing update information per device type (e.g., Apple Watch update support duration).

#### IP CAMERAS

Out of the 50 IP cameras analysed – 25 with update information on coolblue.nl and 25 without – 21 devices (42%) had update information on the manufacturer’s website. Among the 25 devices with update information on coolblue.nl, 15 also had it on the manufacturer’s site. In contrast, only 6 of the 25 devices without update information on coolblue.nl had it on the manufacturer’s website. Notably, for only 3 of the 15 devices where both sources provided update information, the details were consistent. For one of these, an Ezviz camera, coolblue.nl states *24 months after release date* while the Ezviz website states *at least 2(TWO) years after the first shipment for sale [...] we may support for 3 years or longer if serious vulnerability is discovered*.’ This discrepancy matters

because consumers relying solely on coolblue.nl's information would not know that the manufacturer allows extended support in critical situations, which could influence purchase decisions. Similarly, in 11 of the 12 cases where information was inconsistent, the manufacturer listed a longer duration. In these cases, coolblue.nl described the update period as "X months after the release date," while the manufacturer used phrasing such as "at least Y months" after the release date or after the device was last sold—where Y was consistently greater than X. This under-reporting suggests that consumers relying only on retailer information may make suboptimal choices, either avoiding products with longer support or underestimating the manufacturer's broader security commitment.

In our sample of 50 devices, we included devices from all 10 manufacturers due to our stratified sampling strategy, with each manufacturer contributing between one and 13 devices. Only TP-Link had update information on its website for half of its devices (3 of 6), while the other nine manufacturers either consistently provided or did not provide update information for all their devices. Although we usually found update information on the product listing page, for about a quarter of the devices it was listed on a separate webpage, either as a generic update policy for all the manufacturer's IP cameras or as a list of models with their update support durations.

## 5

#### SMART PRINTERS

Only 12 of 50 (24%) smart printers in our sample of 50 devices, had update support information – 8 of which also had update information on bol.com and 4 of which did not. Only two of the 24 brands in our sample provide update information – HP and Oki. HP states '*... products are generally provided with a period of support for firmware updates for 36 months from date of purchase and sometimes longer*' for all of its devices while on Bol, the specified update duration varies across the seven HP printers listed. Only in one case does bol.com also mention an update duration of 36 months, for the majority (5 of 7) it states an update duration of 24 months, while for the remaining one it states 60 months. We observe a similar discrepancy between update duration stated on bol.com for the Oki printer, bol.com states 36 months, while the website states '*Our devices will be supported with security updates at least during the defined support period set out in the statement of compliance of such devices or longer where required by law. In addition, our devices may be supported with updates for up to five (5) years after the product end-of-life.*'

For both HP and Oki, the update information was on a separate page and not with the product listing on the manufacturer site. This implies that only consumers who are keen to double check the update duration will end up identifying the update duration stated by the manufacturers, the others will be influenced by the, often times, lower duration on bol.com or by the lack of update information on the manufacturer's product page.

#### SMART SPEAKERS

Only 9 of 50 smart speakers (18%), from 5 of 24 manufacturers, in our sample have update support information listed on the manufacturer website, six of which also had update information on Bol. One manufacturer, Sony, had the update information for only one of the two speakers in our sample. Of the nine smart speakers with update information, only one had this information directly in the product listing; in the remaining eight cases, it was provided separately. Here as well, we observe discrepancies between the duration specified on bol.com and the manufacturer website. For instance, Google states

*‘Google Nest connected home devices will receive automatic security updates for at least five years from the date we start selling them on the US Google Store’* but the duration on bol.com ranges between 12 and 36 months after the introduction date. Similarly, for a Sony speakers, bol.com states update support till January 2023 but the website promises support till December 2024, a full two years longer than the duration on Bol. In the case of Loewe, the update support duration stated on both are the same, 2 years, but Loewe’s websites has a crucial additional piece of information – *‘In addition we offer [...] latest available security update for that firmware free of charge for a period of 8 years after [...] model is released on the market.’* In all these cases, a consumer only looking at the bol.com website will get an impression that the manufacturer supports with updates for a shorter duration than what the manufacturer states.

### SMART WATCHES

For smart watches, we observe an even split: half of the 50 devices in our sample (50%) have update information available on the manufacturer’s website, while the other half do not. Among the 25 watches with update information on coolblue.nl, only 10 also have this information on the manufacturer’s website. In contrast, 15 watches that lack update information on coolblue.nl do have it listed by the manufacturer. Of the 10 watches with update information available from both sources, six provide consistent details, while for three, coolblue.nl lists a longer update duration, and for one, a shorter duration. A similar even split is observed at the manufacturer level: five out of ten manufacturers provide update information for all their smart watches on their websites, while the remaining five do not. For 16 of the 25 watches with update information on the manufacturer’s website, the information is included directly in the product listing, for the other 9 it is listed on a separate webpage.

Overall, we find that the manufacturer websites contain update information in between 18% (for smart speakers) and 50% (for smart watches) of devices. We also observe a pattern of update duration stated on online stores being lesser than that on the manufacturer websites across all the device types. This again challenges the notion of retailers being passive channels that publish manufacturer information. The retailer’s promise might have more legal meaning than the manufacturer, since the consumer is entering a contract with the retailer, not with the manufacturer. All in all, the analysis for these three device types questions the usability of this information for a consumer.

### 5.6.3. COMPARISON OF RETAILERS, MANUFACTURERS AND DATABASE FOR SMART TVs

Similar to the manual validation of manufacturer websites for the other device types, we also checked manufacturer websites for update information related to smart TVs. There were 47 smart TV manufacturers with update information on bol.com and 11 without. We therefore checked the update information on the manufacturer websites for 72 smart TVs, 47 with update info on Bol, one from each manufacturer and 25 random devices without update info.

We found that only 25 of the 72 TVs (34.7%) had update information on the manufacturer website. The update information was either presented in the product sheet available on the website (in 13 cases) or via a QR that could be scanned to reach the EPREL database

(8 cases) or both (4 cases). This lack of consistency in how the information is presented makes the information less useful from a consumer perspective who now has to spend additional time in figuring out where this information is located.

Of the 47 of 72 devices (65.3%) without any update information on the manufacturer website, three devices have the product information sheet but the update information field is empty. Two other devices have a QR code to take a user to the EPREL website which says the information is not available for that specific device.

With regards to manufacturers, three of the 27 manufacturers have the update information for some TV models but not others, seven have the information for all their models in our sample while the remaining 17 do not have the information for any of their TVs. In some cases these were smaller brands for which we were not able to identify any dedicated manufacturer website or information.

In summary, we find that despite the Energy regulation, manufacturer websites provide access to update information for TVs only in 34.7%. In contrast, the update information was present on manufacturer websites for 42% of IP cameras. As the regulation doesn't require this information to be published online, we next examine the EPREL database to see if it is provided there.

## 5

#### VALIDATION OF CENTRALISED DATABASE

The product information is uploaded by manufacturers and importers into a central European Product Registry for Energy Labelling (EPREL) database<sup>2</sup>, a mandatory step to comply with EU energy labelling regulations. Once the products are registered, the product information sheet can be downloaded in all EU languages. Consumers can access the information for a specific model by searching by brand and model name or by registration number. A QR code linking to the product information sheet on this database is also printed on the energy labels of the products. The ecodesign regulation 2019/2021 also mandates a minimum period of software support: it requires that manufacturers make the 'latest available version' of firmware available for at least 8 years after the last unit of a model is placed on the market, and likewise make security updates available for 8 years, free of charge [4].

We manually checked the EPREL website for the 42 devices without product information sheet or update information on the manufacturer website. We find that 34 of these have the update information on the EPREL database. This indicates that for 20 devices across 15 manufacturers, the information is available in the centralised database but the manufacturers do not link to it in any way from their own website. Moreover, 8 TVs do not have the update information even on the database despite the Energy regulation requiring manufacturers and sellers to provide this information to be able to sell their products [3].

Moreover, for the eight TVs with update information also available on Bol, we find inconsistent information between the bol.com website and the product information sheet from the EPREL database. The duration stated on bol.com ranges from one year to five years, in contrast to the 8 years that is stated in the product information sheet of all the TVs. In one notable case, for the same Samsung TV model, bol.com lists an update duration of 'at least 2 years', the product information sheet says 8 years, while Samsung's

<sup>2</sup><https://eprel.ec.europa.eu/screen/home>

website states at least 5 years. Strictly speaking these are not inconsistent with each other, but in terms of guiding consumer purchasing decisions and providing assurances about updates, the differences are problematic.

## 5.7. DISCUSSION

In this work, we analyse how online stores in the EU disclose update support durations for IoT devices – an important point of contact, since consumers often rely on information on product pages when making purchase decisions. For our first research question on disclosure, we find striking differences across store types. Temu, the Chinese e-commerce platform, provides no update information for any of the device types we examined. By contrast, Amazon and some local EU retailers do disclose update durations, though inconsistently.

To answer our second research question, we quantify these differences. Amazon provides update information for only up to 0.4% of products, despite having a designated field on every product page. Dutch stores show far higher coverage: on bol.com, between 33% and 64% of devices include update durations, while on coolblue.nl the range is even wider, between 15.2% and 91.5%, depending on device type (subsection 5.6.1). For smart TVs specifically, update information is more frequently disclosed: we find such details on local retailers in 4 of 22 EEA countries, and across six European Amazon country sites (Table 5.4).

Finally, to address our third research question, we look beyond retailers and compare disclosures from online stores, manufacturer websites, and the EPREL database for smart TVs. Here, we uncover notable inconsistencies: the same device often carries different stated durations depending on the source, raising concerns about the reliability of the information that consumers ultimately rely on. In line with earlier research [106, 182], we find low levels of update information on manufacturer websites. Despite various regulations placing the responsibility for providing update details on manufacturers [66, 74], update information is available on manufacturer websites for only 18% to 50% of the IoT device types analysed. On a more positive note, due to the Energy Regulation [3], 81% of the smart TVs in our sample include update information in the centralized EPREL database, alongside energy-related details.

**Conflicting Minimum Update Durations** Surprisingly, we observed inconsistencies in the stated minimum update durations across the two Dutch online stores, manufacturer websites, and the centralized database for identical device models. The first inconsistency is simply that one source might have information, while the others do not. This means the information might exist, but it does not propagate through the relevant channels. The second inconsistency is that different sources give different assurances as to the update duration. The two Dutch online store retailers often list different durations for the same product. Which ones should consumers trust? Manufacturer websites do not offer clarity either, as they frequently provide no information or, when they do, different – and typically longer – minimum update durations than those listed by retailers.

We should note that these discrepancies are not strictly contradictory. For instance, when a store states ‘at least two years’, this is technically consistent with an ‘at least five

years' claim from the manufacturer. For a consumer contemplating a purchasing decision, however, they provide quite different signals. They suggest a pattern of systematic under-reporting by retailers. Notably, the retailer's statement might have more legal force than the manufacturer, since the consumer is entering a contract with the retailer, not with the manufacturer. Under most consumer protection laws, particularly within the European Union, it is the retailer who is liable for ensuring that the goods sold are as described and fit for purpose – under the concept of 'product conformity'. Therefore, if a retailer advertises a specific update support duration – regardless of what the manufacturer states elsewhere – consumers may be entitled to rely on that promise as part of the purchase contract.

This raises the possibility that the under reporting by online retailers is not accidental but may reflect a deliberate effort to limit their liability. As stated in [subsection 5.6.2](#), in most cases the update duration stated on the online stores is consistently lesser than that on the manufacturer website. Only in one case was the update duration higher on the retailer website. This suggests that, while retailers rely on manufacturers to provide updates, they are not merely passive conduits of manufacturer information; instead, they may actively determine the minimum update durations they advertise based on internal assessments and their own risk calculus.

## 5

**Policy Interventions** Our findings suggest that while policy interventions can improve the availability of update support information, not all interventions have observable effects. The UK PSTI Act, which legally requires manufacturers to disclose update details to consumers, appears to have had little effect, based on our analysis of [amazon.co.uk](#) and manufacturer websites. In contrast, the Dutch ACM intervention [26] and the EU Energy Labelling Regulation for smart TVs [3] show clearer signs of measurable impact, leading to improved availability of update information. Still, the results are far from ideal: information is often missing, retailer compliance is inconsistent, manufacturers frequently fail to disclose details, and the centralized database contains discrepancies compared to what consumers see elsewhere. Interestingly, we find little evidence of positive spillover effects beyond the jurisdictions directly targeted by regulation. One might expect that once update information is compiled and included in product descriptions, they propagate to other countries and store fronts, but this is rarely the case. For instance, although Amazon has the infrastructure to display update support information across all its EU store fronts, these fields are often left empty and are not carried over even when available on another national site (e.g., [amazon.nl](#) vs. [amazon.de](#)). Similarly, while the centralized database for smart TVs feeds into some local stores, it remains absent from most local online stores and entirely missing on Temu.

**Recommendations** Our analysis shows that regulations are not a silver bullet. While introducing regulation is essential for addressing market failures, reducing information asymmetry [18], and helping consumers make security-conscious purchase decisions, we also observe that compliance and enforcement remain major challenges. The Dutch case, where active steps were taken to verify compliance [164], underscores the importance of enforcement to make regulations more than a paper reality. Dutch stores offer update information in the highest percentage of all the online stores. In parallel with regulation,

there is a need to support the development of systems and tools – such as APIs or indexable widgets [60] – that make it easier for platforms to present update information consistently. These tools could also enable third parties (e.g., NGOs or consumer groups) to aggregate and monitor update data across brands, making it easier for consumers to compare update durations across retailers and manufacturers [71] potentially incentivizing them to compete on update duration.

Regarding information display, we observe that manufacturers either present update information on product pages or in centralized databases — each mode serving a different consumer need. At the purchase stage, showing update details directly on product listings helps inform decisions. Post-purchase, a centralized repository offers convenience by consolidating update information across devices of a manufacturer. Ideally, manufacturers should adopt both approaches: maintaining a centralized database and linking it from individual product pages to maximize transparency and utility.

**Limitations** The data from websites and the centralised database was collected between Dec 2024 and May 2025. The information presented might have changed since then. We suggest future work can do a follow up study after the CRA to understand its effects.

While we observe conflicting values in some cases between the online store, manufacturer website and the centralised database, there is no ground truth. It is tempting to treat the manufacturer disclosure as ground truth, but the store enter into a legal contract with the buyers, so their disclosure actually has immediate legal impact.

We are analyzing claims about duration, not actual duration. It will take years to see if stores or manufacturers live up to their promises. It is beyond the scope of this paper to verify which duration will be followed in practice.

In each online store and manufacturer website, we manually checked the pages for update support information to the best of our ability. We also identified update information that was hidden under energy ratings, available only when scanning a QR code. While there might be corner cases that we might have missed, we are confident that we covered all areas of the stores where a consumer would look for the update information.

## 5.8. CONCLUSION

In this work, we analysed the availability of update duration information on online stores across EEA countries. We find that despite policy efforts, update information is often missing or inconsistent across retailers and manufacturers. Our findings highlight the need for stronger compliance checks and systems that make update information publicly accessible and consistent.



# 6

## VICTIMISATION PATTERN OF IOT BOTNET ATTACKS

*The number of Distributed Denial of Service attacks is growing, and the attack vectors are also changing. The advent of IoT botnets like Mirai has challenged the popularity of older techniques like amplification attacks. In this paper, we characterise the consequences of this change for the victimisation pattern of DDoS attacks. We conduct the first empirical comparison of victims of amplification attacks and botnet attacks and draw on the properties of targets outlined in Routine Activity Theory (RAT) to characterise the differences. We analyse the differences in the victimisation patterns at the level of IP addresses and Autonomous Systems. We find differences in the types of networks the victims reside in; botnet attacks are more common against hosting ASes, and its victims are significantly more likely to use dedicated hosting. We also observe that victims of botnet-based attacks tend to be in ASes with larger customer cone sizes. We use a balanced random forest classifier to distinguish the features of the victims of each attack. The model's output confirms our findings and draws out additional geographical differences in the victim distribution. Using the target properties of Value, Inertia, Visibility and Accessibility outlined in RAT, we find that victims of botnet attacks tend to have higher value and visibility and lower inertia than those of amplification attacks. We explain the differences in these patterns of victimisation to the underlying differences in the attack technique. We further use RAT to outline the policy implications of our analysis.*

### 6.1. INTRODUCTION

Any online service faces the threat of being taken offline due to Distributed Denial of Service (DDoS) attacks. These attacks overwhelm the target with spurious requests, exhaust its server capacity and render the service unavailable. The emergence of the DDoS-as-a-service economy made these attacks more accessible and decreased the barriers to entry [116]. Also termed booters, these services enable customers to purchase and launch DDoS attacks against a target of their choice. The downtime due to these attacks imposes

significant economic costs in terms of availability, recovery and reputation [19]. It has been estimated that an hour of downtime causes between 61K and 67K USD in loss of revenue [193]. Moreover, in some cases, consumers tended to diversify their providers as a countermeasure against DDoS in the aftermath of DDoS attacks [11]. There has also been a significant negative impact on the stock prices of companies when DDoS attacks disrupt the service to consumers [8].

As is often the case [18], the persistence and growth of DDoS attacks can be better explained by a lack of incentives rather than a lack of technical solutions. The actors best positioned to implement controls for DDoS attacks – the intermediaries – do not have sufficient incentives to do so. For instance, Source Address Validation (SAV) is a technical solution that prevents source address spoofing in a network and therefore limits the prevalence of certain types of DDoS attacks. In an empirical study of 334 ISPs, Lone et al. [144] found that 73% of the ISPs have not fully deployed SAV in their networks. They concluded that the lack of complete adoption is due to a lack of incentives for the network operators. While the cost of the implementation is borne by the operator, the benefits are reaped by the rest of the internet.

This lack of incentives is despite most victims of DDoS attacks residing in broadband access networks, with relatively fewer targets in hosting and enterprise networks [169]. However, this victim distribution is for DDoS attacks that use amplification techniques: sending small spoofed packets that trigger an amplifier – such as a DNS, NTP or SNMP server – to send a much larger response to a victim. These attacks are typically purchased from booters which offer low-powered but accessible and cheap attacks to consumers who use them predominantly against other consumers, such as in the context of gaming. While these low-powered attacks place some additional stress on the broadband networks, the consequences are not severe enough for network operators to take stronger preventive actions like SAV.

However, the last few years have seen the emergence of powerful DDoS attacks from IoT botnets. An attacker infects and takes control of a large number of IoT devices to create a botnet. The collective power of all the devices can then be used to launch attacks with magnitudes higher than earlier techniques. The notorious IoT botnet-powered DDoS attack on ‘Dyn’ in 2016 was the largest attack seen till then and disrupted services like Reddit, Twitter and CNN. More recently, in August 2022, Imperva reported on a DDoS attack using IoT botnets that had a total of 25.3 billion requests setting a new record for the largest DDoS attack mitigated by them [84]. Moreover, IoT botnet-based DDoS attacks are not only more powerful, but their proliferation is also increasing at an alarming pace. There are reports that current geopolitical tensions and hacktivism have triggered an increase in the proliferation of botnets [165].

What remains unknown, however, is the consequence of this change on the victimisation pattern. Understanding the impact of the change on the victim distribution is important because it might trigger a subsequent change in incentives. Our current knowledge of victim distribution is shaped by high-profile attacks that make the news and industry reports based on limited visibility from their customer networks. On the one hand, we could argue that the change in the attack vector would not make any difference. The operators of the attack infrastructures are not the actors ordering the attacks; their clients are. Thus, one hypothesis would be that the infrastructure is simply a tool, and the

actor ordering the attack does not care which tool is used as long as it gets the job done.

On the other hand, IoT botnets enhance the magnitude of DDoS attacks and, at the same time, undermine our current DDoS mitigation techniques, like scrubbing [168]. So another hypothesis would be that this combination of increased attacker power and decreased defence capability would enable the attackers to go after better-defended targets, albeit at a higher cost, thus changing the victimisation patterns. Without a sufficient investigation into the victimisation patterns of IoT botnet-based DDoS attacks, we cannot confirm if the change in attack technique has caused a change in the corresponding targets. Thus, it is essential to study the change in victimisation patterns not only because it is under-explored in literature but also because the findings can help us understand the change in and distribution of the incentives. This would be a necessary step in identifying the changes necessary to law and public policy to better align the incentives.

In this paper, we address this gap. We find out what the change in attack vector means for the victims and how this might reshape incentives to invest in DDoS defence measures. We identify the victims of IoT botnet-based DDoS attacks and compare them to earlier attacks using amplification techniques. We do not know if the victim distribution identified in 2015 by Noorizan et al. [169], still holds both for the amplification attacks since then or for the more recent attack vector based on IoT botnets. Thus, our primary research question is, *'Who are the victims of DDoS attacks using IoT botnets, and how does the victimisation pattern of IoT botnets compare to that of earlier attack techniques?'*

We answer the question using two existing data sources on DDoS attacks that are, as yet, under-utilised for studying victimisation patterns. First, we collected attack commands sent by the Command and Control servers (C2s) to Mirai bots from Netlab<sup>1</sup>. Next, for the benchmark, we collected amplification attack data; victim IP addresses from amplifier honeypots dubbed AmpPots [130]. Using these data sets, we compare DDoS commands sent by IoT botnets to honeypot data over 15 months (January 2020 to March 2021). We compare network-level features of the target IP addresses, like the type of Autonomous Systems, and host-level features, like the density of domains hosted on the address, to identify victimisation patterns. Using AmpPot data between January 2016 to March 2021, we also study the longitudinal evolution of victims of amplification attacks over the four years. Further, we map the identified features to the four tenets of Value, Inertia, Visibility and Accessibility outlined in Routine Activity Theory (RAT) and use the framework to evaluate how incentives might play out in terms of suitable targets and defences. In short, our contributions are as follows:

- We perform the first empirical study outlining the change in the victimisation pattern of DDoS attacks due to IoT botnets. We compare and quantify the differences in victims of IoT botnet-based DDoS attacks and amplification attacks.
- We identify victimisation patterns of DDoS attacks by characterising the networks where the victims reside, i.e., network type, ranking, geo-location and network size. Our results show that botnet attacks are proportionally more common against Hosting ASes. 36.6% of botnet attacks were against hosting ASes compared to only 21.1% of amplification attacks. We also find that the victims of botnet attacks are more likely to use dedicated hosting.

---

<sup>1</sup><https://netlab.360.com>

- We also identify the sectors where the targeted victims operate. Victims of IoT botnets are primarily Small and Medium Enterprises. In contrast, online gaming remains the most targeted industry for amplification attacks, accounting for more than a third of the attacks.
- We use a balanced random tree classifier to distinguish the characteristics of victims suffering an IoT botnet attack vs an amplification attack. We identify statistically significant differences in the rankings of the networks where the victims reside – botnet attacks target high-ranked ASes proportionally more. The classifier also outlined geographical differences in the victim distribution. We find that a larger percentage of victims IPs of botnet attacks were in Europe and Africa, while amplification attacks were more prevalent in the Americas and Asia.
- We characterise the differences in the victimisation pattern using the target properties outlined in RAT. We connect the differences identified to the underlying differences between the attack vectors and draw out policy implications.

## 6.2. BACKGROUND AND RELATED WORK

DDoS attacks pose a relevant and significant threat in our current digital landscape. In 2020, DDoS attacks grew more than 50% increasing both in complexity and attack volume [214]. The explosion in network traffic due to the changes caused by the COVID-19 pandemic made it easier for attackers to launch DDoS attacks. Since the servers were already under stress due to high traffic volume, it took a relatively lesser effort to overwhelm the servers with requests and take the service offline. According to an industry report [228], in pure numbers, 25% of all attacks in 2020 were targeted at the technology sector, but the corresponding attack size was relatively low. The healthcare sector, on the other hand, suffered the most in terms of average attack size but was amongst the least attacked industry. 2022 and 2021 saw a slight reduction in the percentage of DDoS attacks since 2020, 9.7% and 3.5%, respectively, but the peak bandwidths in 2021 were almost seven times higher than 2020 [98].

### THE MARKET FOR DDoS ATTACKS

An important contributing factor to the high volume of DDoS attacks is the low entry barriers to launch one. For the tech-savvy and motivated attackers, there are YouTube tutorials on creating botnets and launching DDoS attacks. For the non-tech-savvy, amplification attacks can be purchased online from the aforementioned DDoS booter services with an ease similar to online shopping. In 2020, these cost a mere \$48 for an hour, \$134 for a day and \$1,000 for a month [214]. Booters also have more in common with e-commerce websites beyond ease of purchase. Musotto and Wall [160] showed that booters are similar to e-commerce websites in terms of how their products, price and customers are differentiated and also noted that the profit margins are not very high.

On the defence side, the threat of DDoS attacks has created a market for DDoS Protection Services. There is a prominent trend towards increased adoption of these services, especially by large web hosters [113]. Such protection is typically classified into proactive and reactive protection [108]. Proactive protection is always on, looking out for potential

attacks and, depending on the exact configuration, includes varying levels of packet analysis to determine which packets to block. Reactive protection, on the other hand, only analyses meta-data of network traffic to detect anomalies and traffic mitigation kicks in only when the analysis points to suspicious activity.

### TECHNICAL STUDIES ON DDoS ATTACK: AMPLIFICATION AND BOTNET BASED

Technical studies on understanding amplification DDoS attacks have outlined detection mechanisms [48, 115, 226] while other studies investigate the various protocols that are commonly abused for amplification attacks. They identify commonly used protocols are the UDP-based NTP, LDAP, OpenVPN, ARMS, Ubiquity Discovery Protocol and the like [125] and also observe that there are over 2.5k DDoS attacks in a single day. Kührer et al. [131] reported on the significant diversity in amplifiers used in DDoS attacks and also estimated that TCP handshakes can be abused to cause up to 20x amplification.

Similarly, several studies have contributed to our technical understanding of IoT botnets and specifically of Mirai. A seven-month retrospective analysis of the Mirai botnet [23] studied its emergence, the evolution of its variants, and the competition for vulnerable hosts. Notably, it also pushed forth the understanding that Mirai marks a significant change in the evolutionary development of botnets, both due to the simplicity of its infection vector and its exponential growth. This provided a wake-up call to prioritise the security of IoT devices to prevent such severe DDoS attacks [124]. To that end, Jerkins [109] catalogued vulnerable IoT devices using the same attack vector as Mirai motivating manufacturers to address their poor security practices. Similarly, Rodríguez et al. [189] identified device types and manufacturers of Mirai-infected IoT devices through Web-UI image scans and banner analysis. With regard to the cleanup of Mirai-infected devices, Cetin et al. [45] show that quarantining and notifying infected customers through the ISP has the maximum impact. 92% of infections were remediated within two weeks, and only 5% were reinfected in five months. Others have also pointed to the significant role of broadband ISPs in combating the spread of IoT botnet infections like Mirai [171].

### VICTIMS OF DDoS ATTACKS

Commercial DDoS protection services claim that any business can be a target for DDoS attacks while available prior research [62, 117] on victims places gaming-related services and end hosts at the forefront. Moreover, targets are typically attacked by different types of attacks, and web servers are targeted most often [112]. Other studies on victims were conducted to better understand attacker motives. Abhishta et al., [10] used RAT to analyse the victim properties of 26 DDoS attack events that made the news. They argue that economic reasons are only one of the possible motives and advise companies to monitor the social, political and cultural dimensions of their environment to have a better understanding of the underlying threats. Another study on attacks on Dutch educational institutions [9] lends evidence to this claim. It found a significant correlation between the academic schedules and the attack patterns leading to the conclusion that the attacks were launched by an actor who would have benefited from the disruption to the educational activity.

However, while there is information on the high-profile attacks that make the news, either owing to the target or the severity, there is scarce info on other attacks. It is

important to note that while high-profile targets might have protection and redundancy in place to mitigate the severity of the attacks, other businesses might not have sufficient protection in place to prevent even less severe attacks. Moreover, there is no work on distinguishing the victims or targets of DDoS attacks via botnets and amplifiers.

### 6.3. DATA SOURCES AND METHODOLOGY

As mentioned in the Introduction, to conduct this victimisation study we use previously collected data sets on DDoS attacks that have been under utilised to study victimisation patterns. To study the victimisation of IoT botnets, we collected Mirai attack data from the Network Security Research Lab NetLab 360's website<sup>2</sup>. For the comparison to earlier attack, we used amplification attack data, collected through AmpPots [130]. To the raw data obtained from these sources, additional data was added to enable meaningful analysis.

As mentioned in the Introduction, we draw upon Routine Activity Theory (RAT) to analyse the patterns of victimisation observed in both the types of attacks. RAT posits that crime happens at the convergence of space and time where a motivated offender and a suitable target are present in the absence of a capable guardian [51]. Although originally developed for offline crime, RAT has been adapted to the context of online crime [232]. RAT outlines four properties that affect the target suitability – Value, Inertia, Visibility and Accessibility, often referred to as VIVA. Value is the gains to the attacker from the attack, Inertia is the target's resistance to the attack, Visibility is the degree of exposure of the target to the attacker and Accessibility is the reachability of the target. We map each of the victim attributes analysed to the one of target properties outlined in RAT to study the underlying differences that drive target selection. Figure 6.1 shows an overview of the data analysis process including the mapping of the attributes to RAT properties.

## 6

#### 6.3.1. NETLAB DATA

The Mirai botnet attack data used in this analysis is sourced from the Network Security Research Lab, Netlab 360. They used analyser programs to heuristically analyse and extract C2 domains or IP information from samples of the Mirai malware. They then track these C2 servers and publish the command information received from the C2s. A detailed explanation of their methods to extract configuration data, attack methods and dictionaries of usernames and passwords from Mirai samples and to classify and track its many variants is provided in [142]. As part of their OpenData Project, till mid March 2021, they released portions of the Mirai attack data thus captured on their website<sup>3</sup>. We scraped and downloaded this attack data set from their website over the collection period between Jan 2020 and March 2021. We scraped and downloaded this attack data set from their website over the collection period between Jan 2020 and March 2021. The data contains a snapshot of commands sent by C2s to Mirai infected devices and includes the time of the attack, the target IP and port and the duration of the attack. The C2 server IP address though available is obfuscated. We found minor inconsistencies in the data, such as port numbers outside the range of 0 to 65535. However, these are likely artifacts

<sup>2</sup><https://data.netlab.360.com>

<sup>3</sup><https://data.netlab.360.com/mirai-c2/>

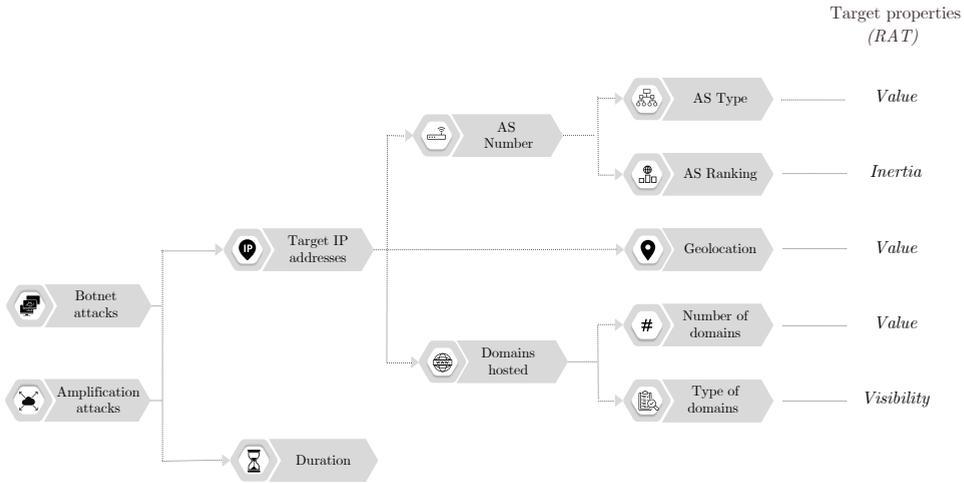


Figure 6.1: Overview of attributes analysed and the corresponding mapping to RAT properties

of Netlab’s data collection methodology or our scraping. Given the low occurrence of such inconsistencies, we omit these data points from our analysis. On average, there were 596 unique target IPs observed each day over the entire period of observation, with a minimum of 11 and a maximum of 1,169 IPs.

### 6.3.2. AMPOT DATA

The Amplification attack data analysed was collected between Jan 2016 and March 2021 through amplifier honeypots, termed AmpPots and its working is explained in detail in the original paper [130]. In short, these honeypots mimic services commonly abused by attackers for amplification attacks and send back legitimate responses. These services include QotD (17/UDP), CharGen (19/UDP), DNS (53/UDP), NTP (123/UTP), SNMP (161/UDP) and SSDP (1900/UDP). Attackers are thus lured into using these honeypots as amplifiers and data on ongoing attacks, targets and techniques are collected by the AmpPots. These amplifiers are deployed in Japan and depending on the ISP their IPs change every 5 to 30 weeks.

Table 6.1: Description of AmpPot data over the period of analysis

Time Period	Number of AmpPots	Types of AmpPots
Jan 2016 to May 2017	9	7 proxied and 2 agonistic
June 2017 to March 2018	7	7 proxied
March 2018 to April 2021	19	11 proxied and 8 agnostic

Over the four years, there were differences in the number and type of sensors used

which are illustrated in Table 6.1. Proxied sensors imitate the functionality of the underlying protocol abused by the attackers. They forward the request to internal servers running the abused protocol and send the responses back to the client. Agnostic sensors, on the other hand, reply with a random bytes of response irrespective of the validity of the request. These operate with the assumption that attackers are more concerned about finding hosts that send back large responses than the validity of those responses. However, in this study, we focus on the larger overarching trends in the victims of these amplification attacks rather than the variations due to the differences in the sensors.

In order to separate attacks from scans, an attack is defined as a series of at least 100 consecutive packets where consecutive is defined as less than 60 seconds apart. This is a change from the 3600s and 600s definition used in the earlier papers [130, 169] but it allows for analysis at a more granular level. On average, the AmpPots observed 9,475 unique target IPs per day over the entire observation period.

The number of attacks per month for each of the data sets is shown in Figure 6.2. The size of the AmpPot data set is higher by two orders of magnitude. However, we are comparing the relative proportions of attacks across various aspects and therefore the difference in absolute sizes does not impact the veracity of the results.

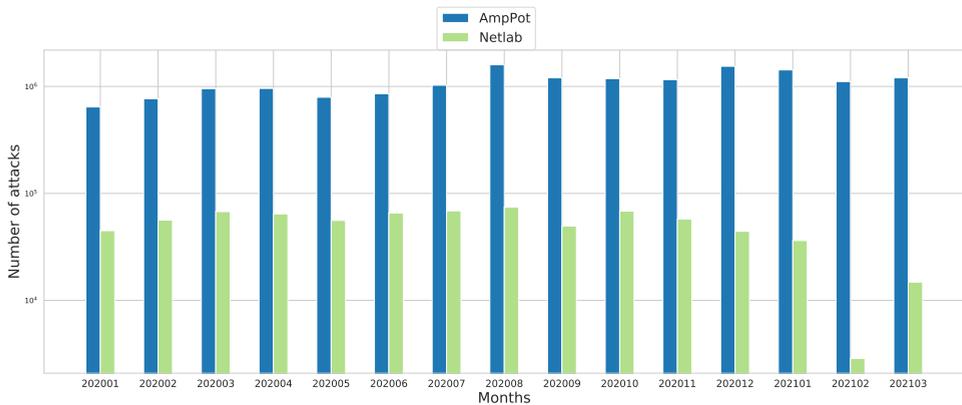


Figure 6.2: Attack distribution over the months

## 6.4. METHODOLOGY

Before describing the methodology used for comparison, we would like to clarify the various terms used in the study. We use the word ‘target’ or ‘target IP’ to denote the entity the attacker intended to affect. However, since we can not directly observe the attacker’s intention, inline with earlier work [169], we use the term ‘victim’ or ‘victim IP’ to refer to the targeted IP address. From both attack data sets we extracted three main attributes for analysis: the duration of the attacks and destination port, both of which were compared directly and the victim IP address, for which additional data was collected for analysis.

### COMPARISON OF AS TYPES.

We first compare the types of Autonomous Systems (ASes) that the victims belong to. Since DDoS attacks also impose significant stress on the networks or ASes that the victims reside in, we refer to these ASes as victim ASes. To get the AS that the target belongs to, we looked up the Autonomous System Number (ASN) of the targeted IP using historical BGP routing data obtained from Routeviews<sup>4</sup>. These files were loaded into the pyASN package<sup>5</sup> freely available for python to perform IP to ASN conversion.

Once we obtained the victim ASN, we used a previously built research database for checking the victim AS type. This database was built manually over the years and is organized around ground truth data from an accurate commercial database - Telegeography GlobalComms Database Service [6]. The mapping accurately distinguishes and labels ASes as broadband ISPs, hosting, governmental, mobile ISP, educational and other types of networks.

In addition, we further improve the classification by identifying hosting ASes using the same heuristic as Noroozian et al. [169]. We classify as hosting any AS that was not classified using the database and contains more than 2,700 second level domains (SLDs). To get the count of SLDs per AS, we use a large passive DNS (pDNS) database provided by Farsight Security [5]. The database contains the mapping of IPs and the corresponding domains that they resolved to over our period of observation. We aggregated the IP level domain count based on the corresponding ASN they belong to and obtained the count of SLDs per ASN.

We corroborated our classification of AS types with those from ASDB [241] which classifies AS types using machine learning techniques on data from RIR's WHOIS and Business Intelligence Databases. We found that both our classification and ASDB have similar coverage rate for our data set. For the Netlab data set, ASDB has 19.6% unknowns while our classification has 12.5%. A similar pattern is observed for AmpPot as well. Since the missing ASes are mostly on the tail-end of the frequency distribution – ASes which are not commonly attacked – they do not have a significant impact on our results.

The main difference between the two databases is in the percentage of ISP broadband and hosting ASes. The percentage of ASes classified as ISP broadband is higher in ASDB compared to our database while the percentage of hosting ASes is lower in ASDB. The difference is due to classification of ISP broadband ASes – ASDB classifies as ISP broadband the ASes our classification marks as Hosting. However, since the ISP broadband ASes in our database have been identified using accurate information from Telegeography and have also been manually validated, we consider our database to be more accurate for our data set. Moreover, we have classified the ASes based on the predominant use of the network, either ISP broadband or hosting, while in some cases other types of users might also be present.

### COMPARISON OF AS RANKING.

To compare the size and connectivity of the victim ASes in the botnet and amplification attack data sets, we used CAIDA's AS Rankings [1]. These rankings are calculated using

---

<sup>4</sup><http://archive.routeviews.org/>

<sup>5</sup><https://github.com/hadiasghari/pyasn>

customer cone sizes derived from BGP routing data and CAIDA's topological data<sup>6</sup>. An AS's rank is inversely proportional to the size of its customer cone – the sum of its direct and indirect customers. The indirect customers are the customers that can be reached through the ASes that a given AS peers with. These ranks denote both the influence of an AS in the global routing system and its size<sup>7</sup>.

### COMPARISON OF VICTIM LOCATION.

We use MaxMind's GeoIP location database<sup>8</sup> to identify the geographical location of the victim IPs in both data sets. The service provides the country an IP address is located in with an accuracy of 99.8%.

### COMPARISON OF DOMAINS HOSTED ON TARGET IPs.

To compare the domain level attributes, we use the same pDNS database described earlier. We retrieved the domains hosted in the top 100 most common IPs for both sets of data across the entire time period of observation. Since we were interested in domains that are hosted, we restricted ourselves to IPs within hosting ASes. We collected the super set of all the domains hosted on these IPs for each month that they were observed in the data sets.

In addition, to get an estimate of the value of the domains hosted in these IPs, we got the Tranco ranking for these domains. Tranco<sup>9</sup> provides a transparent and reproducible popularity ranking of websites. We then manually analysed these domains to identify their types. Since most of the domains did not have an associated Tranco ranking, we analysed the domains with Tranco ranking separately from others without a corresponding ranking.

### MODELLING.

In order to better study the differences in the victimisation patterns of the two attack techniques, we constructed a Balanced Random Forest classifier [47]. Random forest is a supervised machine learning algorithm that uses multiple decision trees to arrive at a final class output. Owing to the differences in the size of our data sets, we used a Balanced Random Forest Classifier that is available as part of the imbalanced-learn library [135]. This uses random under sampling of bootstrapped samples to balance the size.

We only selected properties of victims as features because we were interested in the differences in the victimisation patterns across the two data sets. This implies that the duration and port, although they relate to the victim, were not included in the model since they are properties of the attack itself rather than the victims. There were four features input to the classifier. Of these, two were ordinal – the domain count of the victim IP and the the CAIDA ranking of the victim ASN. The other two features were categorical and were one-hot encoded before being input to the model. These were the region, based on the geo-location of the victim IP, and the AS type of the victim ASN. The countries output by the geo-location were grouped into regions for a more concise representation. The duplicate data points were dropped before being input to the model. We ran the model

<sup>6</sup><https://www.caida.org/projects/ark/>

<sup>7</sup><https://asrank.caida.org/about>

<sup>8</sup><https://www.maxmind.com/en/geoip2-country-database>

<sup>9</sup><https://tranco-list.eu/>

for varying values of number of estimators (the number of trees the model constructs) and maximum depth of the tree and picked the values that had the best accuracy.

## RAT PROPERTIES

As mentioned earlier, we map each of the attribute analysed to one of the four properties of RAT – Value, Inertia, Visibility and Accessibility. We do not use the property of accessibility since all the targets are hosted on the internet and can be reached by any attacker with an internet connection. We mapped the AS level attributes, AS type and ranking, to Value and Inertia respectively. The value gains to an attacker from targeting a specific type of AS, say hosting, will be higher than targeting a broadband customer. The former involves monetary loss from sites hosted on the target IP while the latter will result in lost connection for an individual customer. High ranked ASes will offer higher resistance to the attack and therefore have higher inertia when compared to other low ranked ASes. The location of the target also relates to value. Targets in countries with higher ICT index are more valuable because of the higher dependency of the country on these services. The domain level attributes number of domains and type of domains both relate to visibility. The higher the number of domains that are hosted on a target IP, the higher it's reachability or visibility. Similarly, some types of domains like X are more visible than other types of domains, say Y.

## ETHICAL CONSIDERATIONS.

We adhered to our institution's ethical policy at all times and appropriately handled issues concerning data preservation and data sharing. For the botnet attack data set from Netlab, we notified them about our interest in their data set, and the scraping scripts were designed to minimise the load on their servers. The amplification data set was collected via honeypots. In order for a honey pot to successfully lure attackers, it needs to participate in the attack to a certain degree. However, each honey pot deployment has rate limiting mechanisms to minimise the impact of the participation to a negligible degree.

## 6.5. RESULTS

### 6.5.1. DISTRIBUTION OF TARGETED AS TYPES

We first examine differences between amplification and botnet attack victims by comparing differences among the autonomous systems in which victim IPs reside, i.e. by comparing victim ASes. We compared the distribution of victim AS types identified (as described in [section 6.4](#)) over three dimensions [Figure 6.3](#) – the percentage of unique IPs, the percentage of attacks and the percentage of unique ASes in each data set. The comparison of the distribution of unique IPs across the AS types shows that the highest percentage of victims in both AmpPot (63.8%) and Netlab (47.1%) data sets are in broadband ISPs. Moreover although the second highest category for both is hosting AS, only 14% of AmpPot victims belong to hosting while about 32.4% of Netlab victims are in hosting ASes.

A similar trend is observed in the distribution of unique attacks across the ASes. The most common AS Type for both data sets is ISP broadband (AmpPot - 48.1% and Netlab -

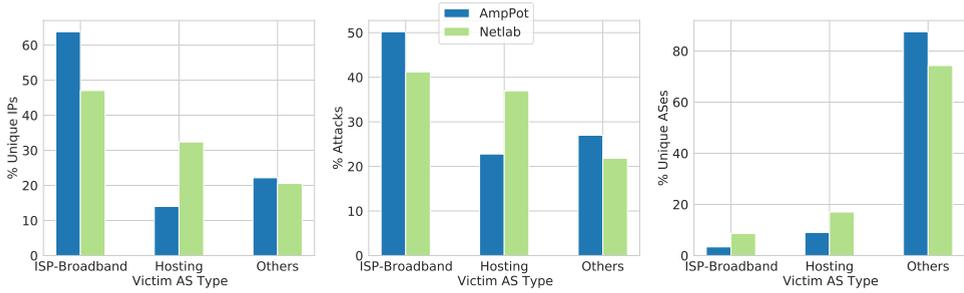


Figure 6.3: Comparison of percentage of unique ASes, attacks and IPs across the AS Types

41.7%). The next highest in AmpPot is Others (30.8%) while for Netlab it is hosting (36.6%). However, when it comes to unique ASes the distribution becomes more interesting. The majority for both (AmpPot 87.6% and Netlab 74.3%) is Others but the second most common is hosting for both (AmpPot 9.4% and Netlab 17.6%) and not broadband ISPs (AmpPot 3.4% and Netlab 8.7%). This shows a remarkable concentration of victims in broadband ASes: 64% of victim IPs in the AmpPot data set are from 3.4% of ASes and 47% of victim IPs in Netlab data set are from 8.7% of ASes. Moreover, even when using ASDB for AS Type classification, the proportional distribution of attacks over the AS types is similar.

Further, we see that the distribution of victim ASes across the AS Types has remained relatively stable when compared to earlier work [169]. The ISP broadband AS type still has the highest number of attacks (48%) and hosts the most unique IPs (64%). However, the percentage of attacks in the Others category which includes education, government, gaming, ISP-Mobile and unknowns among others, has increased.

### 6.5.2. RANKINGS OF THE TARGETED ASes

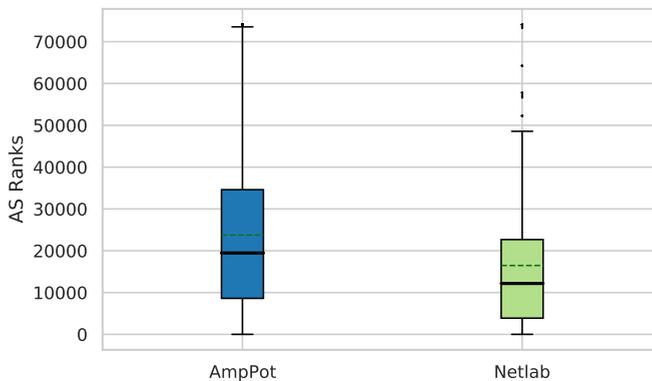
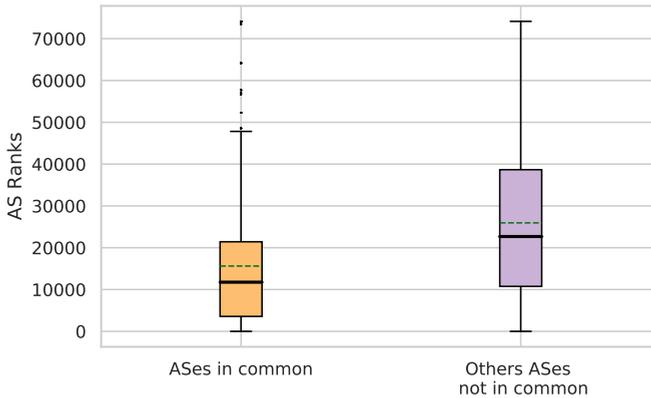


Figure 6.4: Comparison of AS ranks between AmpPot and Netlab data sets

Next we compare the differences in the CAIDA AS rankings of the victim ASes. As



**Figure 6.5:** Comparison of AS ranks between ASes common to both data sets and those unique to each data set

seen in Figure 6.4, we observe a difference in the mean ranking of ASes in the AmpPot and Netlab data sets. The mean rank of ASes in the AmpPot data set is 23,749.9 while for the Netlab data set it is 16,471.5. In order to check if this difference is significant, we performed the Mann-Whitney U test. The results indicate that there is a statistically significant difference ( $p < 0.001$ ). This indicates that the relative size, connectivity, and therefore influence of ASes in the Netlab dataset are higher than those in the AmpPot dataset.

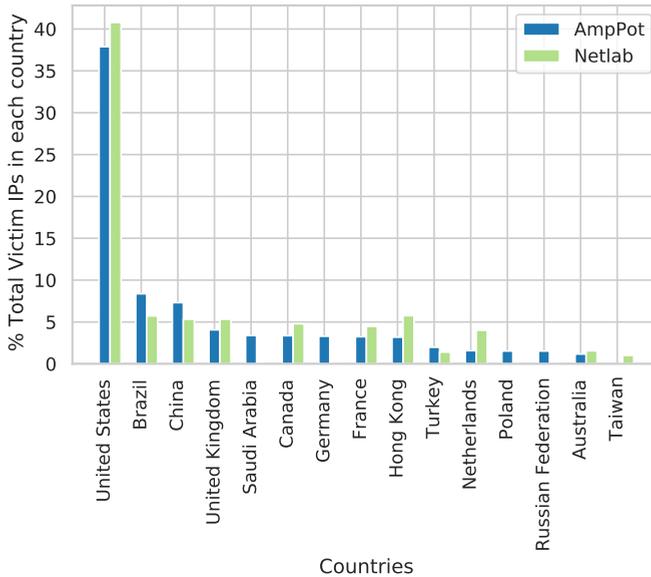
However, it should be noted that this high ranking (and higher influence) is also seen in the ASes that are common to both AmpPot and Netlab as shown in Figure 6.5. The Mann-Whitney U test also showed a statistically significant difference in the ranking of these two sets of ASes ( $p < 0.001$ ). We thus see that the larger, more influential ASes contain victims targeted by both amplification and botnet attacks, but the proportion of botnet attack victims in these ASes are significantly more. 94.7% of unique ASes in the Netlab data set are in this common group compared to 23.6% of unique ASes in the AmpPot data set.

### 6.5.3. GEOGRAPHICAL DISTRIBUTION OF THE TARGETED IP ADDRESSES

We then compared the geographical distribution of victim IP addresses across our datasets. As mentioned earlier, we used MaxMind's GeoIP<sup>10</sup> location to get the geo-location of the victim IPs.

Figure 6.6 shows the top ten countries with the highest percentage of victims. We found that United States has the highest victim IPs in both data sets (Netlab - 44.6% and AmpPot - 37.3%) by a large margin. The next highest country with most victims of botnet attacks is Canada (5.75%) and for amplification it is Brazil (11.6%). Interestingly, in both data sets, China has the third highest number of victim IPs (7.5% in AmpPot and 5.7% in Netlab) and United Kingdom has the fourth highest (3.9% in AmpPot and 5.3% in Netlab). Though Saudi Arabia was the fifth highest country with victim IPs of AmpPot

<sup>10</sup><https://www.maxmind.com/en/geoip2-country-database>



**Figure 6.6:** Countries with more than one percent of victim IPs

## 6

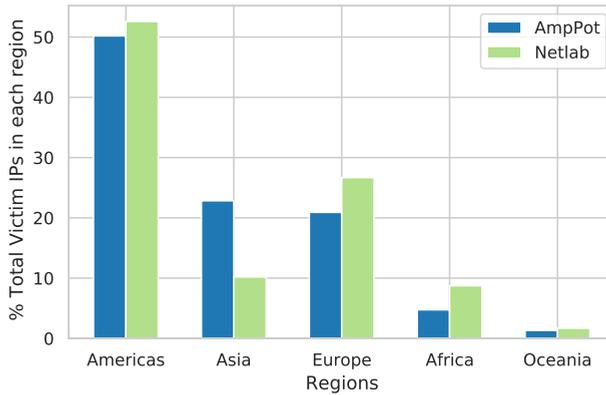
(3.3%), only 0.2% of victim IPs in Netlab were located in Saudi Arabia. The distribution of percentage of attacks across countries with more than one percent of attacks in each data set is shown in Figure 6.6. Although the graph seems to show an over representation of botnet victims in countries with high GDPs, we did not find any statistically significant correlation. The cross comparison of victim location across our amplification attack data and botnet attack data also suggest minor differences.

We then grouped the countries by regions using the Standard area codes provided by the Statistics Division of the UN<sup>11</sup> to study the differences at a more aggregate level. When grouped by regions, the Americas (North America and South America together) rank the highest in both data sets (AmpPot 53.9% and Netlab 52.6%). However, the next highest region with victims of amplification attacks is Asia (21.4%) which has half as many botnet victims (10.2%). The second highest percentage of botnet victims are in Europe (26.7%) which also has 19.4% of amplification victims. We see that the percentage of botnet victims in Africa is twice that of amplification victims. The distribution of the percentage of attacks across all regions is illustrated in Figure 6.7.

### 6.5.4. COMPARISON OF DOMAINS HOSTED IN THE TARGETED IPs

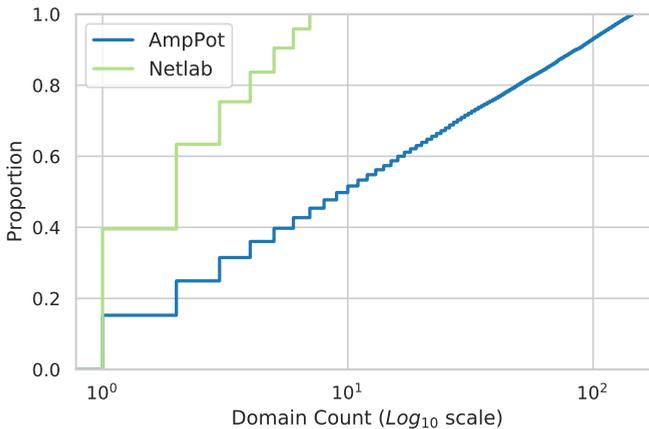
Next, we compare victims by examining domain resources hosted behind the attacked victim IPs in our data sets. As described in the Methodology (section 6.4), we used a passive DNS database to obtain the number of domains hosted on the target IPs in both data sets. We calculated the domain count per IP as the average of the number of domains hosted on the IP through all the months that the IP was seen in the data set. We then

<sup>11</sup><https://unstats.un.org/unsd/methodology/m49/overview>



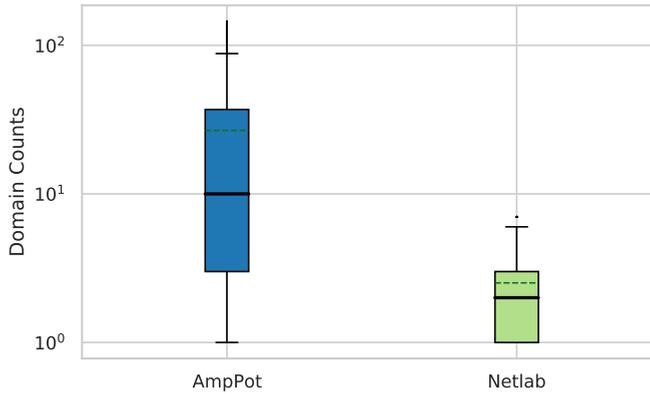
**Figure 6.7:** Percentage of victim IPs across different regions

compared the domain counts of unique attacks on hosting ASes in each data set. We saw that 77% of attacks against hosting ASes in AmpPot and 40% of attacks in Netlab had a domain count of zero. The cumulative distribution function of the remaining attacks with domain counts greater than zero is shown in [Figure 6.8](#).



**Figure 6.8:** CDF of domain count of unique attacks in hosting ASes in both data sets

We observe that both the mean and median domain counts of attacks in the AmpPot data set are higher than that of Netlab. The mean and median domain counts for attacks in the AmpPot data set are 26.7 and 10 respectively, while for Netlab they are 2.5 and 2. The significant difference in domain counts illustrated in the box plot in [Figure 6.9](#) and is also confirmed by the results of the Mann-Whitney U-test ( $p < 0.001$ ). The difference without dropping domain counts of zero is also significant ( $p < 0.001$ ).



**Figure 6.9:** Box plot comparison of domain counts

### 6.5.5. ANALYSIS OF DOMAINS RESOLVING TO TOP 100 MOST COMMON IPs.

In order to compare the types of domains hosted on the victim IPs in each data set, we manually analysed them. We extracted the domains hosted on the top 100 most common IPs in each data sets and dropped outlier IPs with significantly larger domain counts (less than 1% of the IPs). This gave us 274 unique domains on Netlab and 418 on AmpPot.

6

#### COMPARISON OF TRANCO RANKINGS

As mentioned earlier, to get a better estimate of the value of these domains, we got their corresponding Tranco ranking<sup>12</sup>. We observed that 87.6% (240) of domains in Netlab and 96% (401) of domains in AmpPot had no associated Tranco ranking. The domains and corresponding ranking in each data set, where available, are presented in Table 6.2 and Table 6.3. The average Tranco ranking for the domains unique to AmpPot is 254,141.3 while for Netlab it is 199,796.6. If we use the Tranco ranking of popularity as a proxy for value, we see that the targets of AmpPot have lower value than those of Netlab.

#### ANALYSIS OF DOMAINS WITH TRANCO RANKINGS

We analysed the subset of domains with an associated Tranco ranking, separately from the rest of the domains. There were six domains in common between the two data sets. These were two hosting/Cloud providers, three Network Service Company Websites (Norton, RIPE, Geolocation API) and two unreachable domains (KKK.com and KKK.bz). Of the rest, AmpPot has four hosting/Cloud provider sites, two African news sites, one UK LGBTQ website and one unknown (17tahun.com). In Netlab however, the remaining domains have a wider classification: nine hosting/cloud provider sites, four gaming related sites, three domains of messaging platforms (discord, telegram and IRC), two pages linking to drugs, two university websites, one each of a Psychic reading website, a News/Information site, an LGBTQ site and a porn site and finally two unknowns.

<sup>12</sup><https://tranco-list.eu/>

**Table 6.2:** Domains and corresponding Tranco rankings where available - Netlab

Domain name	Tranco ranking
avast.com	725
discord.gg	1242
ovh.com	2561
ripe.net	6674
your-server.de	15249
hetzner.com	25366
hetzner.de	28316
acquia-sites.com	38906
2ksports.com	44098
psychic-readings-for-free.com	46716
zbigz.com	91006
nexus-cdn.com	141504
unsam.edu.ar	167213
dathost.net	179373
dal.net	194209
softether.net	202438
verygames.net	223435
honglingjin.co.uk	251661
aloneproxy.top	259587
sexdrug.tech	343343
fuckarea.biz	345721
bytebx.com	360102
sexwax.me	367564
omgserv.com	392383
lesbian.com	404008
iproxies.club	405449
prick.top	407477
clouvider.net	647979

**Table 6.3:** Domains and corresponding Tranco rankings where available - AmpPot

Domain name	Tranco ranking
secureserver.net	1107
aliyuncs.com	1930
allafrica.com	4292
ripe.net	6674
incapdns.net	8006
transip.net	226059
lgbt.foundation	296282
pro-norton.com	437159
17tahun.com	741612
africanews.org	818292

### ANALYSIS OF DOMAINS WITHOUT TRANCO RANKINGS

Most of the domains in the subset of domains without an associated Tranco ranking did not resolve to an IP address. We therefore ran a check with a domain registration database<sup>13</sup> to get details about the registration status. We found that only 77% and 61 % of domains were still registered as active domains respectively in the AmpPot and Netlab datasets. We took a 100 random samples from each of the domains registered as active for manual

<sup>13</sup><https://who.is/>

analysis. However, despite being registered as active 54% of domains on Netlab and 63% of domains from AmpPot were not accessible. The errors varied from 'Connection refused' (401) to 'Name not resolved'. Of the remaining which were accessible, the distribution is given in Table 6.4.

We see that the most common type of amplification attacks are on gaming related website in line with earlier research [169]. However, although there exist gaming related victims within Netlab, they are not the most popular. The most common victims of botnet attacks are Small and Medium Enterprises (SMEs). Moreover, AmpPot has no domains relating to hosting or Cloud or non-gaming related Servers while Netlab has a fair share of those across both the data sets.

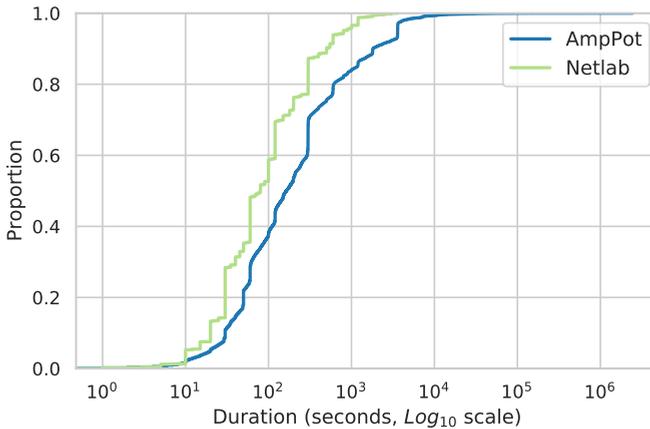
**Table 6.4:** Categories of domains in AmpPot and Netlab

Category	AmpPot	Netlab
Gaming related	26	9
SME	8	18
News/Information	2	-
Network related	1	5
Hosting/Cloud/Server	-	11
Porn/Suspicious	-	3

### 6.5.6. DURATION OF THE DDoS ATTACKS

Next, we examine amplification and botnet attack victims differences by analysing and comparing the duration of attacks directed at each victim across data sets. The average duration of an attack in the AmpPot data set is 754.94 seconds, with a median of 164 seconds. However, while the median duration in Netlab is 80 seconds, the average duration is much higher 793,220.66 seconds (about nine days). The high average is undoubtedly driven by a few outliers in the data. These outliers include values like 4,294,967,295 seconds – the maximum value possible in 32 bits (0xFFFFFFFF) – which amounts to about 136 years. We looked at the Mirai source code<sup>14</sup> and found that the attack function returns an error when the duration value is greater than 3600 seconds. We therefore dropped attacks with duration higher than 3600 seconds for the comparison (about 1.1% of total attacks). Since the duration values in the AmpPot data set are obtained through observation of actual attacks via a honeypot, we did not drop any outliers. The longest attack in the AmpPot data set is 2,466,626 seconds (about 28.5 days).

<sup>14</sup><https://github.com/jgamblin/Mirai-Source-Code/blob/master/mirai/cnc/attack.go>



**Figure 6.10:** CDF of duration of attacks

Figure 6.10 shows the cumulative distribution function for the duration of both data sets. After dropping the outlier durations in Netlab, the average duration is 195.5 seconds and the median is 80 seconds. The results of the Mann-Whitney U-test shows that the differences in the duration are significant ( $p < 0.001$ ).

### 6.5.7. MODELLING

As outlined in the Methodology (Subsection 6.4), we ran a Balanced Random Classifier model to check for differences in the features of the victim in the two sets. We got the highest accuracy (0.66) with a maximum depth of 8 and 500 trees; the feature weights output by the model are shown in Figure 6.11.

The results show that highest contributor to the difference between the victims is the domain count of the IPs followed by the ranking of ASes. This is also in line with our analysis which shows significant differences in both the domain counts and the AS ranking. The differences across regions, though lesser by an order of magnitude, is mostly similar to our region analysis. However, where our analysis only showed minute differences in the percentage of victims in the region of Americas, we see that it is the second highest driving factor for differences in region. Upon closer inspection into this divergence using a visual tree interpreter tool, we found that the model uses a higher value for the Americas region as a higher weightage for the AmpPot class. Interestingly, there are only negligible differences amongst the contribution of the different AS types to the classification.

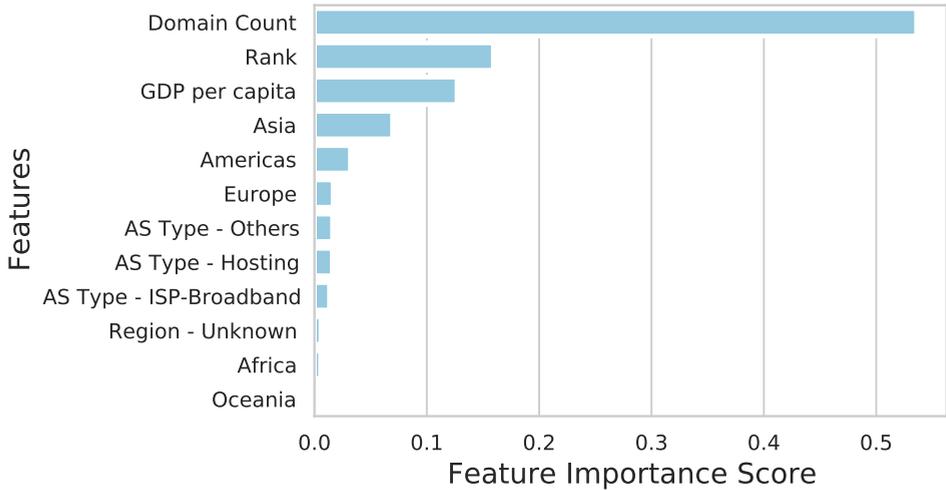


Figure 6.11: Random Forest feature importance scores

## 6

### 6.6. DISCUSSION

Our primary aim with this research was to identify the victimisation pattern for IoT botnets and compare it to the pattern for amplification attacks. We postulated that if there is an underlying difference in the actors deploying these attacks and the corresponding targets, we would see this difference reflected in the victimisation pattern. Conversely, if the attack vector is irrelevant to the actors ordering the attacks, we would not see any difference in the victimisation pattern.

Our results show clear differences in the victimisation pattern within all six attributes analysed. The AS type analysis indicates that a larger percentage of botnet victims reside in hosting ASes. The CAIDA AS ranking shows that relatively more botnet victims are in high-ranking ASes that serve a larger consumer base. The geographical differences show a higher percentage of botnet attacks in countries with higher GDP per capita, albeit with a few exceptions. At the IP level, we observe that victim IPs of botnet attacks within hosting ASes have a higher domain density and also host more popular domains, as indicated by the available Tranco rankings. The classifier highlights the significance of domain count and AS ranking in driving the differences in the victimisation pattern and also draws out other geographical differences. Next, we place these results in the context of the RAT dimensions for an accessible target – Value, Inertia and Visibility. We have not considered accessibility since all the targets are accessed via the internet.

#### VALUE

Value refers to the gains for the attacker from attacking the target, monetary or otherwise, like status or prestige. Of the six attributes that we analysed, three relate to value – AS type, geolocation and the number of domains. For all three, we find that the higher value of a target correlates with a higher prevalence of botnet attacks on the target. We postulate that higher the value of the target, higher the value to the attacker from attacking the

target.

As mentioned earlier, hosting ASes have a higher value from an attacker's perspective than broadband ASes. Hosting ASes charge higher for their service compared to broadband ASes. Therefore, a deterioration in the quality of service due to a DDoS attack on one client will have a higher impact on the revenue for hosting ASes than broadband ASes. Our results show that though broadband ASes are the most common target for both types of attacks, botnet attacks are relatively more common against hosting ASes.

Similarly, countries with higher GDP per capita and ICT development index have a higher dependency on IT services and therefore derive a higher value from them. Although not statistically significant, at a country level, we find a higher occurrence of botnet attacks against victims in the US, UK, Canada, Germany, France and the Netherlands.

Finally, the number of domains hosted on an IP is a clear indicator of value due to the difference in pricing structure between shared hosting and dedicated hosting. Shared hosting, with a larger number of domains per IP, is cost-effective and easy to use, while dedicated hosting with fewer domains per IP offers a more stable and predictable performance. This enhanced performance comes at a higher cost and also requires expertise to set up and maintain. Therefore, dedicated hosting is apt for high-value domains with higher traffic and bandwidth requirements, while shared hosting is an attractive option for personal websites and domains with less traffic. While there is no hard threshold for what counts as shared hosting, prior studies have put it at around ten domains per IP [208, 209]. We find that the average number of domains on victim IPs is 10 for botnet attacks and 26 for amplification attacks. Thus, by this metric, AmpPot victim domains are more likely to use shared hosting services, while botnet victims are more likely to have dedicated hosting.

### INERTIA

Inertia refers to the resistance offered by the target to the attacker. This could relate to the size of the targets, attacks on smaller targets are easier to execute than those on larger targets, or the defence capability of the target. In this research, we mapped the CAIDA ranking of the ASes to inertia. We find that lower inertia or higher resistance correlates with a higher percentage of botnet attacks.

Botnet attacks are more prevalent against high-ranking ASes. These high-ranked ASes have larger customer cone sizes and higher revenue compared to lower-ranked ASes. DDoS attacks on these ASes will impose a higher societal cost since a larger number of prefixes are reachable through these ASes. These ASes thus have both the means – due to their higher revenue – and the motive – to decrease the impact of attacks – to invest in DDoS protection.

### VISIBILITY

Visibility refers to the degree of exposure of the target to the attacker. Our manual analysis and classification of domains resulted in six types: Gaming related, SME websites, News/Information, Network related, Hosting/Cloud/Server and Porn/Suspicious. Of these, we group gaming-related, network related and hosting/cloud/server as low visibility domains because the exact domains are known only to those who have intimate knowledge of these services. For instance, the domain to access the configuration of a

hosting or cloud service includes the public domain of the service, say 'aliyuncs.com', but has additional sub-domains like 'susharefile.oss-cn-shenzhen.aliyuncs.com'.

We find that these low visibility domains are more common in amplification attack victims, while high visibility domains like SMEs are more common in botnet attacks with one exception. Two domains related to News/Information are among the amplification attack victims, while there are no domains related to news/information within the botnet attack victims.

The framework of RAT thus helps us understand that there are differences in targets or victims driven by the differences in attack type. Botnet attacks are more common against high-value victims with a lower inertia and higher visibility, while amplification attacks are more common against low-value targets with high inertia and low visibility.

#### SQUARE PEGS AND SQUARE HOLES

We see that due to the differences in the attack type, each attack type lends itself more suited to a particular target. For instance, botnet attacks are better suited for high-ranked ASes precisely because they might have DDoS prevention and mitigation measures. Evading the defences and launching a successful attack is easier with botnets due to the differences in attack traffic.

Traffic from botnets has legitimate operating system-generated protocol headers, which match the statistical distribution typically observed at the application layer. Due to this similarity with legitimate traffic, machine learning-based mitigation techniques do not achieve high accuracy rates when identifying botnet attack traffic [92, 168]. In addition, the diversity of botnets, each with unique characteristics [188], makes detection more difficult. On the other hand, an attacker exploiting amplification vulnerabilities in a protocol uses specific values in certain header fields to trigger an amplified response. These characteristic header field values make it relatively easy to distinguish an attack from legitimate traffic [28].

Moreover, DDoS protection techniques, like DDoS fingerprinting [99] that propose sharing rules derived from fingerprints of attack sources, can be effective against amplification attacks due to the limited number of amplification sources. However, they will not help defend against botnet attacks since the attack sources are diverse and distributed across the entire IPv4 space. Thus, the added difficulty of evading botnet-based DDoS attacks makes them a more attractive option against high-value clients.

The costs of each of the techniques might also explain the higher prominence of botnet attacks against high-value targets. Amplification attack infrastructures are easier to maintain than botnets. The public services that can be abused for amplification are widely available; the attackers only need to cover the cost of scanning for open amplifiers and launching attacks on demand.

On the other hand, IoT botnets need constant renewal since infections of most IoT malware families are non-persistent; a power cycle removes the infection. Even if the bot remains infected, the connection with the C2 servers is often lost within a few days, as C2 addresses are hard-coded in the binaries. So, as soon as the C2 is taken down, the bots are stranded [210]. This makes the upkeep of an IoT botnet onerous and time-consuming, potentially driving up its operators' costs. In contrast, amplification services are widely available and often masquerade as benign stressor websites, and they are not subject to similar take-down efforts.

While we could not find any reliable data specifying the costs for each of the attack techniques, the above-mentioned factors support the conjecture that botnet attacks are priced higher. The superior attack power, higher operational costs and the resultant lower availability might play a role in price differentiation between amplification attacks and botnet attacks. The higher price of botnet attacks would also deter the less motivated attackers, e.g., teenagers wanting to win a game of Minecraft, since cheaper options are available. This, in turn, can also explain the higher prominence of high-value domains in the botnet attack data set. Thus, like square pegs matching square holes, certain types of attacks match best with a certain type of target.

#### IMPLICATIONS FOR LAW AND POLICY.

The 2016 study on victims of amplification attacks [169] found that the low price of the attacks attracted behaviour that, like vandalism and file sharing, is strictly speaking illegal but not a profit-driven crime. On the other hand, in our results, we see that botnet-based attacks are more costly to execute and go after more valuable targets than, say, individual gamers and thus causing more economic losses. These are more likely part of a profit-driven cybercrime operation or political action. From a crime prevention and mitigation perspective, the policies for these two types of problems are very different.

For low-level crimes like vandalism and file sharing among consumers, the government usually pursues strategies like awareness campaigns and administrative law mechanisms, like statutory fines. In the past, these strategies have been used to tackle the consumer demand for DDoS attacks. An awareness campaign launched by the UK's anti-cybercrime agency warned youth who searched for DDoS booter services online about the illegality of DDoS attacks [175]. This proved to be effective at keeping new users out of the DDoS markets [52].

For profit-driven crime, on the other hand, they usually approach it via criminal law and law enforcement actions. As botnet techniques further evolve and become more widespread, we can expect to see a higher proportion of these attacks in the DDoS landscape. Thus, more of the mitigation will fall to criminal investigations and disruption efforts rather than consumer-focused interventions.

The framework of RAT further highlights the two main possibilities for the prevention of DDoS attacks. We can decrease the number of suitable targets by limiting the attack power of the IoT botnets. This implies having stronger security measures in our IoT devices and thereby decreasing the size of the botnets. Further, we can also increase the capability of the guardians by improving our ability to defend against IoT botnets.

This can be best done via information sharing, which has shown to be effective in mitigating cybercrime [154]. An anti-DDoS coalition, NoMoreDDoS, established in the Netherlands, enables information sharing and collaboration amongst partners to collectively tackle the threat of DDoS attacks [22]. The partners include government organisations, internet service providers and internet exchanges, among others. This illustrates that network intermediaries like high-ranked ASes, the most common targets of botnet attacks, are uniquely positioned to benefit from and initiate such information sharing. They have the incentive to share information – to minimise the stress on their networks. Further, the size of these ASes would also protect them against a potential negative backlash of being cut off from their peers.

## 6.7. CONCLUSION

Denial of Service attacks are almost as old as the Internet, yet our analyses prove that they show no signs of disappearing. In fact, they are growing in number and diversifying in terms of attack vectors. With the increase of tools to identify services that can be misused to perform amplification attacks, DDoS attacks have been commodified and are accessible to criminals with all kinds of different skills. On top of these amplification services, the appearance of IoT botnets made possible to control millions of devices which in turn are also being used to also launch DDoS attacks. All this together has attracted a great variety of criminals launching DDoS attacks against a wide range of victims.

While this research does not demonstrate a displacement of traditional booter services, the commoditization of IoT botnets may alter the market's technological supply. We did not observe repeated victimisation to increase over the years, meaning that the increase in frequency of DDoS attacks is tied to a growth on the number of victims. However, we see that the newer targets operate in large variety of sectors, with small medium enterprises getting more attacks and gaming services still being a common target.

In the arms race between attackers and defenders, the defence measures must routinely adapt to attack techniques to minimise the impact of an attack. As we observe with our results, the newer techniques, which are tougher to defend against, attract enterprising attackers motivated to cause damage to high-value clients. So, amplification and botnet attacks are not of the same feather and their victims are not flocked together. We see that like square pegs finding square holes, victims who are better able to defend against DDoS attacks are attacked using more robust techniques. This stresses the urgency to adopt stronger legal actions against miscreants launching DDoS attacks.

# 7

## CONCLUSION

This dissertation analysed the presence, type and context of market signals for IoT S&P on e-commerce platforms. We have presented five studies in Chapters 2 - 6, that together help unpack different dimensions of the market signals for IoT S&P like consumer perception, manufacturer practices, sellers and platform design and broader societal consequences. Our main research question was:

*What signals for security and privacy are present in the e-commerce platforms that sell IoT devices?*

In this chapter we first provide a brief summary of the findings from the five studies presented earlier. We then discuss the broader societal implications of these findings in [section 7.1](#), while also focusing on how the organically occurring signals on e-commerce platforms can be strengthened to support the different actors in making more S&P-aware decisions. In [section 7.2](#) we outline the governance implications and in [section 7.3](#), we suggest directions for future work.

### CHAPTER 2 – S&P SIGNALS WITHIN CONSUMER REVIEWS

In this chapter, we investigated to what extent customer reviews of IoT products provide S&P information to consumers at the point of purchase. We found that around 10% of consumer reviews contains S&P signals which included technical statements about features, frustrations with specific device use activities, as well as vignettes about trying to use a device in a particular context. Negative views on IoT devices were reflected in generally lower overall ratings for devices. All in all, we found that customer reviews provide a valuable and widely-used mechanism for conveying S&P information to consumers – prior to, and complementary with, potential future labelling schemes for IoT. We recommended mechanisms to amplify these signals within the e-commerce platforms to encourage consumers to making more S&P conscious purchase decisions.

### CHAPTER 3 – IOT DEVICE SALES AND S&P SIGNALS

In Chapter 3, we analysed whether consumer preferences for S&P are reflected in their real life purchase decisions through studying the sales of IoT devices and the signals for S&P in the e-commerce platforms. We constructed a model with the device sales as the independent variable and price, expert S&P rating and overall rating, user ratings and review count from two e-commerce platforms as the predictor variables. Our results showed that despite lack of information about S&P for a majority of the devices, a one standard deviation increase in the S&P rating of a device was associated with a 56% increase in sales. However, we also observed that the relationship was moderated by the price of the device. The effect was stronger at lower prices and decreased corresponding to an increase in price.

In addition, we found that availability of complete update support information on online stores reflected the update practices of manufacturers and can therefore act as a signal for S&P. Further, we also find that on one of the online stores, devices with complete update support information correspond to a 69% increase in sales. This suggests that there are positive incentives at play that will reward manufacturers and sellers who adopt S&P transparency initiatives like security labels with higher sales. Our results also highlighted the crucial role of sellers in ensuring the success of such initiatives since they are responsible to update the relevant information on online stores.

### CHAPTER 4 – THE EVOLUTION OF IOT S&P

We presented our study on the evolution of S&P features in IoT devices over the past decade in Chapter 4. We studied 428 IoT devices from 23 manufacturers, focusing on three widely used and mature IoT device types: IP cameras, smart printers, and smart speakers. Our findings showed that while IP cameras maintained consistently high S&P ratings, smart printers exhibited lower ratings with a slight declining trend, and smart speakers had the lowest ratings with no clear temporal pattern. At the manufacturer level, only a minority (3 out of 23) demonstrated improvement, while the majority (12 out of 23) maintained stable S&P ratings, and the remaining eight showed no clear trend in the ratings. Our analysis also uncovered a surprising trend of inconsistent deployment of S&P features in subsequent device models of the same manufacturer. We find that stability in ratings obscured such underlying inconsistencies. This highlights a need to help manufacturers operationalise S&P best practices within their development processes to ensure a more systematic and coherent focus on S&P across the development process.

### CHAPTER 5 – IOT UPDATE INFORMATION ON ONLINE STORES

In Chapter 5, we examined the availability of update support information for IoT devices on e-commerce platforms in the EEA countries. We analysed over 26,000 product listings across major online retailers, including regional EU stores, Amazon, and Temu, as well as manufacturer websites and – for smart TVs – the centralised EU EPREL database. We identified significant gaps and inconsistencies in the information available to consumers. However, our findings also show that regulatory measures and targeted transparency interventions are associated with better information disclosure. Dutch e-commerce platforms, which were subject to a policy intervention for update support disclosure, had high rates of availability, while platforms like Temu and Amazon disclosed little to

none. The centralised EU energy labelling database, which contains update support information alongside energy ratings for electronic displays such as smart TVs, had the highest percentage of availability, likely due to stronger enforcement mechanisms. We also observed discrepancies in the update duration across different sources raising concerns about consistency and accuracy. For the same device, the update duration on the e-commerce platform differed from that on the manufacturer websites and, for some smart TVs, also from the duration stated in the centralised energy database. Our findings highlight both the potential and the limitations of current regulatory efforts, and underscore the need for mechanisms to ensure not just disclosure, but also accuracy.

## CHAPTER 6 – NEGATIVE EXTERNALITIES FROM IOT BOTNETS

This chapter focused on understanding where the negative externalities from lack of strong S&P measures in IoT devices lie, specifically in the context of IoT botnets. Our research showed that the commoditization of IoT botnets may alter the technological supply in the DDoS booter services. We did not observe repeated victimisation to increase over the years, meaning that the increase in frequency of DDoS attacks is tied to a growth on the number of victims. However, we see that the newer targets operate in large variety of sectors, with small medium enterprises getting more attacks and gaming services still being a common target. This illustrates the broader negative externalities of DDoS attacks stresses the urgency to adopt stronger legal actions against miscreants launching DDoS attacks.

### 7.1. DISCUSSION

In this dissertation, we analysed various market signals for IoT S&P with the aim of better understanding the extent to which S&P related information is available on e-commerce platforms. The broader thrust for this research is the idea that e-commerce platforms are not merely passive conduits that connect consumers to sellers and manufacturers. They actively shape purchase outcomes through their rating, review and recommender systems [191]. With this perspective, it becomes important to better understand the interplay between S&P signals embedded in the e-commerce platforms and their influence on IoT purchase decisions. We now reflect on the findings from the five studies conducted as part of this dissertation through the roles of each stakeholder followed by a discussion on the broader societal implications.

#### 7.1.1. CONSUMERS

As we discussed in the Introduction, consumers are one of the key players in the online market for IoT devices. While an individual consumer might not have much power to influence the S&P choices of manufacturers, the aggregated decisions of all consumers in the market is collectively reflected in the market demand and has a tremendous influence on manufacturers. It is therefore essential to understand this collective market demand for IoT S&P or lack thereof.

In Chapter 2 we identified that on average across the six Amazon web stores considered, around 9.8% of consumer reviews mention S&P concerns organically. We also identified seven themes that these concerns articulate. Our finding complements other

studies that find that consumers care for S&P. It is important to note that a consumers concerns about S&P might be different from that of other actors like researchers or consumer interest groups. For instance, while multi-factor authentication is generally considered a good security measure, we find some reviews expressing friction with it because it delays access to crucial notifications.

Although many e-commerce platforms provide mechanisms for consumers to contact sellers directly, this does not necessarily guarantee that S&P-related concerns would be relayed to the actual manufacturers, especially considering cases where third-party sellers are involved or when brand-licensing relationships obscure accountability. As such, there is value in exploring more formalised and transparent channels that ensure consumer feedback, especially on S&P issues, reaches manufacturers and helps them improve their product. We already noticed how some reviews, especially negative ones, elicited replies from sellers who offered solutions, replacements or at the least assured improvements in subsequent releases. This suggests that consumer reviews already function, to some extent, as a feedback loop. However, making this channel more visible and systematised could enhance its effectiveness. A more transparent review-response mechanism – where manufacturers visibly engage with consumer S&P concerns – could also foster greater competition among manufacturers to address such concerns proactively, thereby enhancing consumer trust.

Further, our findings from Chapter 3 throw light on the relationship between S&P, consumer demand and price. Our results put two common narratives around IoT S&P and price in context – that consumers prioritise price over S&P and that consumers are willing to pay more for S&P. Our results show that consumers are willing to pay more for S&P up to a certain price point beyond which this preference decreases. Consumers are therefore sending a strong signal about their preference for IoT S&P, however this preference is also tempered by their spending constraints. While the exact causal mechanism that links better S&P with higher sales is unclear, it is nonetheless important to note that the market rewards devices with better S&P with higher sales – despite the information asymmetry in the market. If additional information is provided to help consumers make more S&P informed purchase decisions, there might be a more acute increase in the relationship between S&P and sales. However, our overall analysis of the e-commerce platforms highlights how the current design neither provides consumers the information needed to make S&P conscious purchase decisions nor does it allow for easy comparison between products based on S&P.

Despite regulations insisting that the update durations of IoT devices need to be displayed on the website, a majority of the products do not have the information. Moreover, there are other design elements that can be considered to facilitate more S&P informed purchase decisions. For instance, these platforms can choose to highlight S&P related reviews. As we argued in Chapter 2, irrespective of whether these reviews are correct, misinformed or fake, there is value in highlighting them because it will nudge consumers into considering S&P aspects of IoT devices prior to purchase. Similarly, an additional rating specifically for the S&P aspects of the devices can be introduced. Such mechanisms might push manufacturers into providing official S&P information to balance the user perspectives of S&P. Moreover, it will also encourage them to make the S&P features more user friendly lest consumers give it a low rating. It is important to note, however, that there

is currently no incentive for the e-commerce platforms to take these measures unless there is a policy intervention. We revisit this in the governance mechanisms outlined below.

### 7.1.2. MANUFACTURERS

Many of the findings from this research are particularly relevant to manufacturers. From the organic emergence of consumer concerns and demand for S&P outlined in Chapter 2, to the limited availability of update information on online stores discussed in Chapter 5, the central role of manufacturers becomes evident. The positive correlation between stronger S&P features and higher sales identified in Chapter 3 is a strong signal for consumer demand for IoT S&P which can serve as an incentive for manufacturers to prioritise S&P.

Our findings in Chapter 4 show that although the S&P ratings have remained broadly stable over the years, there are inconsistencies in the implementation of S&P features across successive IoT device models from the same manufacturer. A "glass half full" perspective of the result would be that, despite the variation, the stability of the ratings suggests that a baseline level of S&P is still maintained. Conversely, the "glass half empty" perspective of the inconsistent implementation would be that the stability obscures the lack coherent processes to address S&P within the development lifecycles. This underscores the need for manufacturers to focus not only on S&P outcomes but also on the internal processes that drive them. There is ample research in organisational behaviour that highlights the importance of organisational processes and also provides an insight into how such efficient processes can be designed. There is need for more interdisciplinary and transdisciplinary research on how insights from organisational psychology can help manufacturers design more efficient processes to improve the S&P not just for IoT devices but also across other areas. Our results from Chapter 3 which show that between two IoT devices of the same price, a consumer will choose the one with better S&P features can serve as an incentive for manufacturers to make cohesive S&P policies. The results also imply that manufacturers who can lower the cost of implementing S&P can gain a competitive advantage further highlighting the need to streamline the S&P development processes.

Another area that manufacturers can streamline on to improve their S&P signalling is the communication of S&P features to consumers. The lack of information on update duration even on manufacturers websites again points to a systematic lack of integration of S&P into the product life cycle. While it is understandable that disconnects exist between teams responsible for deciding on the update support duration and those managing online content and communication, we found that many other product features are clearly communicated on these platforms. This further highlights the importance of process-level improvements — such as incorporating update duration into the standard set of product information disseminated from the product teams to the teams who post them online.

### 7.1.3. E-COMMERCE PLATFORMS

Our research showed the presence of signals for S&P within the e-commerce platforms and provided insights on how these signals can be used to better support S&P-conscious

purchase decisions. As we discussed earlier (in subsection 7.1.1), there can be mechanisms to highlight reviews that mention S&P and provide a S&P specific user rating. In addition, we also note that the update information on e-commerce stores, even where available, is often towards to bottom of the page along with other details like product dimensions and weight. These might not be checked often, especially for IoT devices. It would therefore be beneficial to have a separate section to present the S&P related information. Nonetheless, the question remains: what incentives do these platforms have to implement such mechanisms?

While there might not be immediate or obvious business incentives, we do find instances where platforms take proactive steps to support informed consumer choices. For example, Coolblue.nl offers a dedicated section on cybersecurity for internet connected devices. Their product selection guides often include S&P considerations<sup>1</sup>, including the update support duration. This can help increase consumer awareness about S&P and help them make more informed purchase decisions. From coolblue.nl's perspective, presenting this information can have strategic, reputational, and practical advantages. Such transparency initiatives bolster its positioning as a consumer-oriented retailers and enables increased trust which can help increase customer loyalty and brand reputation. Moreover, such guides are also SEO-optimised. At the time of writing this dissertation, the coolblue.nl's page ranks among the top 10 in the google search results (from a Dutch IP) for "smart home online privacy". This implies that any consumer looking for S&P information might visit the coolblue.nl's page and end up purchasing from them. This shows that there are positive incentives for platforms to proactively implement mechanisms to aid more S&P aware purchases. This is crucial since the platform design has a huge potential to influence consumer purchase decisions.

We already observe some areas where platforms might be playing a more active role in the S&P of the devices they sell – in specifying the update support durations. We found a pattern of under reporting update durations on e-commerce platforms when compared to manufacturer websites (in cases where such information was available). As we discussed in our paper, this may be a deliberate strategy by platforms to limit their own liability, since the purchase contract is typically between the consumer and the platform, not the manufacturer. Alternatively, platforms may estimate the update duration based on prior knowledge of manufacturer practices and opting for a more conservative value as a precautionary measure. However, this also presents an opportunity for manufacturers and e-commerce platform providers to engage in dialogue aimed at establishing mechanisms that facilitate clear and transparent communication of update support duration information from manufacturers to consumers. For both e-commerce platforms and the manufacturers, providing accurate, consistent information about the update support durations will serve to increase their reputation and signal a stronger degree of commitment to consumers and aid their decision making.

#### 7.1.4. SOCIETAL IMPLICATIONS

We studied the consequence of the advent of the new types of DDoS attacks powered by IoT botnets to the society through the changes in the victimisation patterns of IoT botnets. We observed that these new techniques are used on better protected targets

<sup>1</sup><https://www.coolblue.nl/en/advice/smart-security-that-is-safe-for-your-online-privacy.html>

thereby creating higher strain on societal resources. This highlights the criticality of using all possible avenues to improve the S&P of IoT devices rather than merely restricting to technical counter measures.

For instance, one of the other avenues could be improving consumer awareness about S&P of IoT devices. One of the reviews we analysed – for the study presented in Chapter 2 – explicitly mentioned that the router does not come with a unique default password and urged users to change the password as soon as possible. Interestingly, this is exact attack vector that is leveraged by the Mirai botnet to create large hordes of IoT botnets. This shows that IoT vulnerabilities do permeate the society and are voiced, unprompted, by consumers in online reviews. As mentioned earlier, it has not only been established that consumers read consumer reviews prior to purchase but also that purchase moments are opportune moments of change which such initiatives can capitalise on. As we discussed in the study, while it is important the companies and organisations who are the victims of botnet attacks invest in measures to protect themselves, it is also important that we mitigate the creation of such botnets in the first place. Empowering consumers to make more S&P informed purchase decisions is a crucial step in that direction.

## 7.2. IMPLICATIONS FOR GOVERNANCE AND POLICY MAKING

We now take a step back and address the insights from the dissertation that can help strengthen the governance and policy making around IoT S&P. Governance refers to the totality of interactions in which the government, other public bodies and civil societies participate to solve societal problems. There are three main mechanisms through which the interactions between these different stake holders take place – hierarchical, network and market [151]. Each of these mechanisms has a distinct feature. Hierarchical transactions are top-down, centralised and led by the authority of one actor. Market transactions are decentralised and characterized by the supply demand (dis)equilibrium with coordination determined by prices. Network transactions are peer-to-peer with the authority distributed among mutual dependencies among the stakeholders. Although the focus of this dissertation is on the market mechanisms that promote IoT S&P, there is undoubtedly a mix of all three mechanisms in the governance landscape of IoT S&P. In this section, we outline the aspects of IoT S&P governance within each of these mechanisms and provide suggestions for how they can be optimised to improve the S&P outcomes.

### 7.2.1. HIERARCHICAL GOVERNANCE

The various legislations in the IoT policy landscape are clear instances of hierarchical governance. We find various forms of legislations implemented across different jurisdictions that aim to improve the S&P of IoT devices. The US Cyber Trust Mark [79], the UK Product Security and Telecommunications Infrastructure (PSTI) Act [66], the Cyber Resilience Act (CRA) [74] in the EU are also examples of different legislations that put varying degrees of restrictions and requirements on the S&P of IoT devices. This is definitely crucial in improving the S&P of these devices. However, where push meets the shove is in the compliance framework within which the regulations operate. As we saw in Chapter 5, regulations that were implemented under a mandatory compliance regime with strong penalties, as in the case of the mandatory Energy directive or under a compliance mon-

itoring regime like the Dutch intervention have a higher probability of being followed in practice. With the upcoming CRA in the EU, which requires a minimum of 5 years of update support for IoT devices, it is important to keep in mind that without actual follow through on how long a manufacturer actually provides these updates, the 5 year minimum duration risks being mere window dressing. We observed this in the case of smart TVs where manufacturers mentioned the mandatory 8 years of update support in the centralised database but presented a lower update duration on their websites.

A crucial aspect of legislations is that a manufacturer might have to cater to different, in some cases conflicting, regulations when selling across different countries and/or jurisdictions. This can place significant overhead on manufacturers, especially smaller vendors who might not have sufficient resources to have a legal compliance department that handles such issues. In order to ensure such legislations do not stifle innovation, it is important for countries to enter into agreements to recognise each others legislations, compliances or certifications. For instance, if a manufacturer can prove compliance to a given legislations, which can be enabled through third party certification authorities, those certificates can be recognised in other countries. Similarly, countries can establish arrangements to recognise each other's cybersecurity labels. The international agreement between Finland and Singapore is a an example of how such arrangements can work in practice [80]. Such arrangements harmonise the requirements across the respective labels, allowing products that meet the security criteria to be placed on the market in both countries.

### 7.2.2. NETWORK GOVERNANCE

A network governance model is characterised by mutual dependencies. The dominance of a single stakeholder is replaced by mutual dependencies among all actors, creating a networked, collaborative form of governance. One contributing factor to the poor implementation of S&P features in IoT devices is the distributed nature of the IoT supply chain, which creates a vacuum of responsibility — each actor assumes that S&P will be addressed either upstream or downstream [158]. In Chapter 4, we argued that the inconsistent implementation of S&P features across different device models may also reflect underlying changes in vendors. Some vendors may provide native support for certain features within their firmware, while others may not, leading to variability in S&P.

Within this fragmented supply chain, there is potential for stakeholders to self organise around IoT S&P and collaborate to improve it. Especially in light of stricter regulatory requirements, such collaborations can help prevent each vendor from having to 'reinvent the wheel' when implementing S&P. Instead, they can strategically direct their efforts thereby increasing overall efficiency and consistency across the ecosystem.

Moreover, operationalising a regulation like the CRA typically requires horizontal coordination among various stakeholders, including manufacturers, standardisation bodies, cybersecurity agencies, civil society organisations, and third party certification providers. For instance, the International Organisation for Standardisation (ISO) has a special working group for the CRA including various stakeholders like manufacturers, market surveillance authorities, policymakers and consumer organisations that represent consumer interest. Such forms of network governance create a forum that serves to ensure each stakeholder's perspectives are fully considered in translating a regulation

to a standard, ultimately improving the uptake of the standard. However, currently, e-commerce platforms are not part of these forums. Given the crucial role they play in ensuring the transparency of S&P information, especially through its placement on their platforms, other mechanisms could be explored to help them understand the importance of such S&P initiatives. The intervention by the ACM in the Netherlands, which specifically engaged with e-commerce platforms, shows one possible direction such initiatives can take.

### 7.2.3. MARKET GOVERNANCE

The main focus on this dissertation was on the signals for S&P in the market for IoT devices. Our argument was that if there are strong signals for IoT S&P present in the market, they can decrease the information asymmetry and serve as incentives for manufacturers to increase their focus on S&P. As mentioned earlier, we do find evidence of such signals both organically like in customer reviews and through interventions like the display of update support durations. However we also note that in their current forms, these signals are scattered and cannot reliably serve to decrease the information asymmetry significantly. More focused effort is needed for these signals to collectively help users in making more S&P informed purchase decisions.

The classic way in which the market can help allocate the resources for S&P more efficiently would be through increased consumer demand. However, in the absence of such strong consumer demand – and the information asymmetry – external interventions like the CRA regulations can help stimulate demand for S&P. Once CRA compliance information is available across all the e-commerce stores, it would nudge consumers to factor S&P into their IoT purchase decisions. Another outcome could be that civil society or consumer organisations could aggregate this information and present it in easy to compare format which can further help consumers make informed purchases. Some organisations, such as Mozilla’s Privacy Not Included and various national consumer bodies, already perform such aggregations and evaluations. This can, in turn, result in manufacturers competing on S&P thereby improving the overall level of S&P in the IoT landscape.

## 7.3. FUTURE WORK

For too long, we have studied security as if it were solely a technical problem (build better defenses), a regulatory problem (mandate better practices), or a user problem (educate better behavior). This dissertation presents the complementary perspective of market signals and design: the presence of S&P information in e-commerce platforms, whether manufacturers are rewarded for security investment, whether consumers can act on their security preferences, and ultimately whether billions of IoT devices flooding into homes worldwide are secure or vulnerable. While our findings offer a better understanding of the current ecosystem and the dynamics at play, they are also several interesting directions that future research can explore.

### LONGITUDINAL ASSESSMENT OF REGULATORY IMPACT

One clear direction for future work is to conduct a longitudinal analysis of the impact of emerging regulations, such as the CRA. It would be interesting to analyse not only

the manufacturer compliance – especially whether they adhere to the stated minimum update support duration of 5 years – but also on how information related to the CRA is communicated in e-commerce platforms. Taking the results presented in this dissertation as a baseline, future research can explore whether the regulation increased the availability of S&P related information on the e-commerce platforms.

#### CONSUMER INTERPRETATION AND DECISION-MAKING

Another area that warrants further investigation is the causal factors underlying the correlation between high S&P ratings and increased sales in our study presented in Chapter 3. While we found strong evidence that better S&P corresponds to higher sales and that the relationship is mediated by price, our study was not set up to investigate the causal mechanisms. This can be addressed in future research, including more empirical examination on how the trade-off between price and S&P manifests in consumer purchase decisions.

#### IMPACT OF E-COMMERCE PLATFORM DESIGN

Another promising direction for further research is to empirically examine the role of e-commerce platforms, and their design, on consumer purchase decisions. Our results have showed that these platforms are not passive intermediaries, but more active gatekeepers of information. For example, future studies could investigate to what extent is the availability and content of S&P information on the e-commerce platforms influenced by liability concerns. It would also be valuable to understand if platform-led S&P interventions like internal rating systems or product guides influence consumer perceptions compared to other sources of S&P information. These would help further unpack the role of market governance in shaping IoT S&P outcomes.

# BIBLIOGRAPHY

- [1] [n. d.]. CAIDA AS Rank. <http://as-rank.caida.org/>. [Accessed: October 10, 2025].
- [2] 2012. *Overview of internet of things (IoT)*. Technical Report Y.4000. International Telecommunication Union, ITU-T Study Group 20. <https://handle.itu.int/11.1002/1000/11559> Approval process: AAP. Status: In force. Former ITU-T Y.2060 renumbered as ITU-T Y.4000 on 2016-02-05 without modification.
- [3] 2019. Commission Delegated Regulation (EU) 2019/2013 of 11 March 2019 supplementing Regulation (EU) 2017/1369 of the European Parliament and of the Council with regard to energy labelling of electronic displays. <https://eur-lex.europa.eu/eli/reg/2019/2013/oj>. Annex V.
- [4] 2019. Commission Regulation (EU) 2019/2021 of 1 October 2019 laying down ecodesign requirements for electronic displays pursuant to Directive 2009/125/EC of the European Parliament and of the Council. <https://eur-lex.europa.eu/eli/reg/2019/2021/oj>. Annex II, E.1.
- [5] 2022. Farsight security. <https://scout.dnsdb.info/dashboard> [Accessed: October 10, 2025].
- [6] 2022. TeleGeography: GlobalComms Database Service. <https://www2.telegeography.com/en/globalcomms-database-service> [Accessed: October 10, 2025].
- [7] Noura Abdi, Kopo M Ramokapane, and Jose M Such. 2019. More than smart speakers: Security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 451–466.
- [8] Abhishta Abhishta, Reinoud Joosten, and Lambert JM Nieuwenhuis. 2017. Analysing the impact of a DDoS attack announcement on victim stock prices. In *2017 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*. IEEE, 354–362.
- [9] Abhishta Abhishta, Marianne Junger, Reinoud Joosten, and Lambert J.M. Nieuwenhuis. 2019. Victim routine influences the number of DDoS attacks: Evidence from Dutch educational network. In *2019 IEEE Security and Privacy Workshops (SPW)*. 242–247. <https://doi.org/10.1109/SPW.2019.00052>
- [10] Abhishta Abhishta, Wouter van Heeswijk, Marianne Junger, Lambert JM Nieuwenhuis, and Reinoud Joosten. 2020. Why would we get attacked? An analysis of attacker’s aims behind DDoS attacks. *Journal of wireless mobile networks, ubiquitous computing, and dependable applications* 11, 2 (2020), 3–22.

- [11] Abhishta Abhishta, Roland van Rijswijk-Deij, and Lambert JM Nieuwenhuis. 2019. Measuring the impact of a successful DDoS attack on the customer behaviour of managed DNS service providers. *ACM SIGCOMM Computer Communication Review* 48, 5 (2019), 70–76.
- [12] ACM (Autoriteit Consument & Markt). 2020. Consumenten beter geïnformeerd over updates bij aankoop slim apparaat na actie acm. <https://www.acm.nl/nl/publicaties/consumenten-beter-geinformeerd-over-updates-bij-aankoop-slim-apparaat-na-actie-acm> [Accessed: October 16, 2023].
- [13] M Billur Akdeniz, Roger J Calantone, and Clay M Voorhees. 2014. Signaling quality: An examination of the effects of marketing-and nonmarketing-controlled signals on perceptions of automotive brand quality. *Journal of Product Innovation Management* 31, 4 (2014), 728–743.
- [14] George A Akerlof. 1978. The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*. Elsevier, 235–251.
- [15] Mitsuaki Akiyama, Shugo Shiraishi, Akifumi Fukumoto, Ryota Yoshimoto, Eitaro Shioji, and Toshihiro Yamauchi. 2023. Seeing is not always believing: Insights on IoT manufacturing from firmware composition analysis and vendor survey. *Computers & Security* 133 (2023), 103389.
- [16] Alina. 2020. 19 zero-day vulnerabilities affect millions of IoT devices worldwide. *Bitdefender HotforSecurity Blog* (22 Jun 2020). <https://www.bitdefender.com/en-us/blog/hotforsecurity/19-zero-day-vulnerabilities-affect-millions-iot-devices-worldwide> Accessed: February 25, 2025.
- [17] Naveen Amblee and Tung Bui. 2011. Harnessing the Influence of Social Proof in Online Shopping: The effect of Electronic Word-of-Mouth on Sales of Digital Microproducts. *International journal of electronic commerce* 16, 2 (2011), 91–114.
- [18] Ross Anderson. 2001. Why information security is hard-an economic perspective. In *Seventeenth Annual Computer Security Applications Conference*. IEEE, 358–365.
- [19] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel JG Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. 2013. Measuring the cost of cyber-crime. *The economics of information security and privacy* (2013), 265–300.
- [20] Ross Anderson, Eireann Leverett, and Richard Clayton. 2019. Standardisation and certification of the ‘Internet of things’. (2019). <https://doi.org/10.17863/cam.35286>
- [21] Ross Anderson and Tyler Moore. 2006. The economics of information security. *science* 314, 5799 (2006), 610–613.
- [22] Anti-DDoS-Coalitie. 2023. About the coalition. <https://www.nomoreddos.org/en/about-the-coalition/> [Accessed: October 10, 2025].

- [23] Manos Antonakakis, Tim April, Michael Bailey, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, Yi Zhou, Manos Antonakakis Tim April, Matthew Bernhard Elie Bursztein, Jaime J Cochran Zakir Durumeric Alex Halderman Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever Zane Ma, Joshua Mason, and Nick Sullivan Kurt Thomas. 2017. Understanding the Mirai Botnet. *USENIX Security '17* (2017). <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis>
- [24] Renana Arizon-Peretz, Irit Hadar, and Gil Luria. 2021. The importance of security is in the eye of the beholder: Cultural, organizational, and personal factors affecting the implementation of security by design. *IEEE Transactions on Software Engineering* 48, 11 (2021), 4433–4446.
- [25] Hala Assal and Sonia Chiasson. 2018. Security in the software development lifecycle. In *Fourteenth symposium on usable privacy and security (SOUPS 2018)*. 281–296.
- [26] Authority for Consumers and Markets (ACM). 2024. Consumers are now better informed about updates when purchasing smart devices, thanks to acm intervention. <https://www.acm.nl/en/publications/consumers-are-now-better-informed-about-updates-when-purchasing-smart-devices-thanks-acm-intervention> Accessed: February 25, 2025.
- [27] Dax K. Basdeo, Ken G. Smith, Curtis M. Grimm, Violina P. Rindova, and Pamela J. Derfus. 2006. The impact of market actions on firm reputation. *Strategic Management Journal* 27, 12 (2006), 1205–1219. <http://www.jstor.org/stable/20142408>
- [28] Yana Bekeneva, Nikolay Shipilov, and Andrey Shorov. 2016. Investigation of protection mechanisms against drdos attacks using a simulation approach. In *Internet of Things, Smart Spaces, and Next Generation Networks and Systems*. Springer, 316–325.
- [29] Giampaolo Bella, Bruce Christianson, and Luca Viganò. 2017. Invisible security. In *Security Protocols XXIV: 24th International Workshop, Brno, Czech Republic, April 7-8, 2016, Revised Selected Papers 24*. Springer, 1–9.
- [30] Christopher Bellman and Paul C van Oorschot. 2020. Best Practices for IoT Security: What Does That Even Mean? *arXiv preprint arXiv:2004.12179* (2020).
- [31] Ketan Bhardwaj, Joaquin Chung Miranda, and Ada Gavrilovska. 2018. *Towards IoT-DDoS Prevention Using Edge Computing*. Technical Report.
- [32] Hugo LJ Bijmans, Tim M Booi, and Christian Doerr. 2019. Just the tip of the iceberg: Internet-scale exploitation of routers for cryptojacking. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 449–464.
- [33] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet Allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.

- [34] 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.
- [35] Farzam Boroomand, Aija Leiponen, and Gurneeta Vasudeva. 2022. *Does the market value attention to data privacy? Evidence from us-listed firms under the gdpr*. Technical Report. Wharton Mack Institute working paper.
- [36] Nellie Bowles. 2018. Thermostats, locks and lights: Digital tools of domestic abuse. *The New York Times* 23 (2018).
- [37] Irina Brass, Leonie Tanczer, Madeline Carr, and Jason Blackstock. 2017. Regulating IoT: Enabling or disabling the capacity of the Internet of Things? *Risk & Regulation* 33 (2017), 12–15.
- [38] Irina Brass, Leonie Tanczer, Madeline Carr, and Jason Blackstock Word. 2017. *Title: "Regulating IoT: Enabling or Disabling the Capacity of the Internet of Things?"*. Technical Report. <https://www.euractiv.com/section/innovation-industry/news/commission-plans->
- [39] Irina Brass, Leonie Maria Tanczer, Madeline Carr, Miles Elsdén, Jason J Blackstock, Jason J. Blackstock, and Jason J Blackstock. 2018. Standardising a moving target: The development and evolution of IoT security standards. *IoT 2018* (2018). <https://doi.org/10.1049/cp.2018.0024>
- [40] 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.
- [41] Mollie E. Brooks, Kasper Kristensen, Koen J. van Benthem, Arni Magnusson, Casper W. Berg, Anders Nielsen, Hans J. Skaug, Martin Maechler, and Benjamin M. Bolker. 2017. glmmTMB Balances Speed and Flexibility Among Packages for Zero-inflated Generalized Linear Mixed Modeling. *The R Journal* 9, 2 (2017), 378–400. <https://doi.org/10.32614/RJ-2017-066>
- [42] Peter Bruce and Andrew Bruce. 2017. *Practical statistics for data scientists*. O'Reilly Media.
- [43] Michel Callon. 2021. *Markets in the making: Rethinking competition, goods, and innovation*. Princeton University Press.
- [44] Fanny Cambier and Ingrid Poncin. 2020. Inferring brand integrity from marketing communications: The effects of brand transparency signals in a consumer empowerment context. *Journal of Business Research* 109 (2020), 260–270.
- [45] Orcun Cetin, Carlos Ganan, 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. Internet Society. <https://doi.org/10.14722/ndss.2019.23438>

- [46] George Chalhoub, Martin J Kraemer, Norbert Nthala, and Ivan Flechais. 2021. “It did not give me an option to decline”: A longitudinal analysis of the user experience of security and privacy in smart home products. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [47] Chao Chen, Andy Liaw, Leo Breiman, et al. 2004. Using random forest to learn imbalanced data. *University of California, Berkeley* 110, 1-12 (2004), 24.
- [48] Chih-Chieh Chen, Yi-Ren Chen, Wei-Chih Lu, Shi-Chun Tsai, and Ming-Chuan Yang. 2017. Detecting amplification attacks with software defined networking. In *2017 IEEE conference on dependable and secure computing*. IEEE, 195–201.
- [49] Judith A Chevalier and Dina Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* 43, 3 (2006), 345–354.
- [50] Victoria Clarke and Virginia Braun. 2014. *Thematic Analysis*. Springer New York, New York, NY, 1947–1952. [https://doi.org/10.1007/978-1-4614-5583-7\\_311](https://doi.org/10.1007/978-1-4614-5583-7_311)
- [51] Lawrence E Cohen and Marcus Felson. 1979. Social change and crime rate trends: A routine activity approach. *American sociological review* (1979), 588–608.
- [52] Ben Collier, Daniel R Thomas, Richard Clayton, and Alice Hutchings. 2019. Booting the booters: Evaluating the effects of police interventions in the market for denial-of-service attacks. In *Proceedings of the internet measurement conference*. 50–64.
- [53] Barry Collins. 2020. Google home hub shows random nest cam footage on family’s device. <https://www.forbes.com/sites/barrycollins/2020/05/15/google-home-hub-shows-random--nest-cam-footage-on-familys-device/> [Accessed: October 16, 2023].
- [54] European Commission. 2014. Radio Equipment Directive (RED). [https://single-market-economy.ec.europa.eu/sectors/electrical-and-electronic-engineering-industries-eei/radio-equipment-directive-red\\_en](https://single-market-economy.ec.europa.eu/sectors/electrical-and-electronic-engineering-industries-eei/radio-equipment-directive-red_en) Accessed: 2024-11-14.
- [55] European Commission. 2022. Cyber Resilience Act. <https://digital-strategy.ec.europa.eu/en/library/cyber-resilience-act> Accessed: 2024-11-14.
- [56] European Commission, Joint Research Centre, E Leverett, R Clayton, R Anderson, and G Baldini. 2017. *Standardisation and certification of safety, security and privacy in the ‘Internet of things’*. Publications Office of the European Union. <https://doi.org/doi/10.2760/47559>
- [57] Federal Trade Commission. 2018. FTC approves final order in asus privacy case. <https://www.ftc.gov/news-events/news/press-releases/2016/07/ftc-approves-final-order-asus-privacy-case> Accessed: 2024-11-14.

- [58] Federal Trade Commission. 2023. FTC Says Ring Employees Illegally Surveilled Customers, Failed to Stop Hackers from Taking Control of Users' Cameras. <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-says-ring-employees-illegally-surveilled-customers-failed-stop-hackers-taking-control-users> Accessed: 2024-11-14.
- [59] Connected Commerce Council. 2023. New research: Online marketplaces help level the playing field for small sellers. <https://connectedcouncil.org/new-research-online-marketplaces-help-level-the-playing-field-for-small-sellers/> Accessed: 2025-05-01.
- [60] Lorrie Faith Cranor, Yuvraj Agarwal, and Pardis Emami-Naeini. 2024. Internet of things security and privacy labels should empower consumers. *Commun. ACM* 67, 3 (2024), 29–31.
- [61] Geng Cui, Hon-Kwong Lui, and Xiaoning Guo. 2012. The effect of online consumer reviews on new product sales. *International Journal of Electronic Commerce* 17, 1 (2012), 39–58.
- [62] Jakub Czyz, Michael Kallitsis, Manaf Gharaibeh, Christos Papadopoulos, Michael Bailey, and Manish Karir. 2014. Taming the 800 pound gorilla: The rise and decline of NTP ddos attacks. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. 435–448.
- [63] 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*. 97–115.
- [64] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. 2018. Breaking! A typology of security and privacy news and how it's shared. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [65] Henry De and Graft Acquah. 2010. Comparison of akaike information criterion (aic) and bayesian information criterion (bic) in selection of an asymmetric price relationship. *Journal of Development and Agricultural Economics* 2 (02 2010), 1–6.
- [66] Department for Science, Innovation and Technology (DSIT). 2023–2024. The uk product security and telecommunications infrastructure (product security) regime. <https://www.gov.uk/government/publications/the-uk-product-security-and-telecommunications-infrastructure-product-security-regime> Accessed: February 25, 2025; Last updated May 2, 2024.
- [67] Hongying Dong, Hao Shu, Vijay Prakash, Yizhe Zhang, Muhammad Talha Paracha, David Choffnes, Santiago Torres-Arias, Danny Yuxing Huang, and Yixin Sun. 2023. Behind the Scenes: Uncovering TLS and Server Certificate Practice of IoT Device Vendors in the Wild. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. 457–477.

- [68] Sujay Dutta and Sandeep Bhowmick. 2009. Consumer responses to offline and online low price signals: The role of cognitive elaboration. *Journal of Business Research* 62, 6 (2009), 629–635.
- [69] 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.
- [70] P. Emami-Naeini, J. Dheenadhayalan, Y. Agarwal, and L.F. Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of IoT Devices?. In *Proceedings of the 32nd USENIX Security Symposium*.
- [71] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2021. An informative security and privacy “nutrition” label for internet of things devices. *IEEE Security & Privacy* 20, 2 (2021), 31–39.
- [72] 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*. 1–12.
- [73] ENISA. 2020. Guidelines for Securing the Internet of Things. <https://www.enisa.europa.eu/publications/guidelines-for-securing-the-internet-of-things> Accessed: 2024-11-14.
- [74] European Commission. 2024. Cyber resilience act. <https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act> Accessed: February 25, 2025.
- [75] European Parliament and Council of the European Union. 2019. Directive (eu) 2019/770 of the european parliament and of the council of 20 may 2019 on certain aspects concerning contracts for the supply of digital content and digital services. , 27 pages. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32019L0770> Accessed: February 25, 2025.
- [76] European Parliament and Council of the European Union. 2019. Directive (eu) 2019/771 of the european parliament and of the council of 20 may 2019 on certain aspects concerning contracts for the sale of goods, amending regulation (eu) 2017/2394 and directive 2009/22/ec, and repealing directive 1999/44/ec. , 28-50 pages. <https://eur-lex.europa.eu/eli/dir/2019/771/oj/eng> Accessed: February 25, 2025.
- [77] Michael Fagan, Katerina N Megias, Karen Scarfone, and Matthew Smith. 2020. *Foundational cybersecurity activities for IoT device manufacturers*. US Department of Commerce, National Institute of Standards and Technology.
- [78] 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>. [Accessed: October 10, 2025].

- [79] Federal Communications Commission (FCC). 2023–2025. Cybertrustmark. <https://www.fcc.gov/CyberTrustMark> Accessed: February 25, 2025.
- [80] Finnish Transport and Communications Agency Traficom. 2021. International cooperation helps finnish businesses enter the market of secure smart devices. <https://www.traficom.fi/en/news/international-cooperation-helps-finnish-businesses-enter-market-secure-smart-devices> Accessed: 2025-05-14.
- [81] Joshua Fogel and Samson Zachariah. 2017. Intentions to use the yelp review website and purchase behavior after reading reviews. *Journal of theoretical and applied electronic commerce research* 12, 1 (2017), 53–67.
- [82] Nathaniel Fruchter and Ilaria Liccardi. 2018. Consumer attitudes towards privacy and security in home assistants. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3170427.3188448>
- [83] FTC. 2015. IoT privacy and security in a connected world. <https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things-privacy/150127iotrpt.pdf>
- [84] Nelli Klepfish Gabi Stapel. 2022. Record 25.3 Billion Request Multiplexing DDoS Attack Mitigated by Imperva. <https://www.imperva.com/blog/record-25-3-billion-request-multiplexing-attack-mitigated-by-imperva/> [Accessed: October 10, 2025].
- [85] Ryan J Gallagher, Kyle Reing, David Kale, and Greg Ver Steeg. 2017. Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics* 5 (2017), 529–542.
- [86] Vaibhav Garg. 2021. A lemon by any other label. *ICISSP* (2021), 558–565.
- [87] 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.
- [88] David Godes and Dina Mayzlin. 2004. Using Online Conversations to Study Word-of-Mouth Communication. *Marketing science* 23, 4 (2004), 545–560.
- [89] S. Gopavaram, J. Dev, S. Das, and L. J. Camp. 2021. IoT marketplace: Willingness-to-pay vs. Willingness-to-accept. In *Proceedings of the 20th Annual Workshop on the Economics of Information Security (WEIS 2021)*.
- [90] Shakthidhar Reddy Gopavaram, Jayati Dev, Sanchari Das, and Jean Camp. 2019. *Iotmarketplace: Informing purchase decisions with risk communication*. Technical Report. Working Paper, 2019, ftp://svn.soic.indiana.edu/pub/techreports/TR742.pdf.

- [91] The Guardian. 2020. Amazon's ring camera comes under lawsuit over hacking threats. <https://www.theguardian.com/technology/2020/dec/23/amazon-ring-camera-hack-lawsuit-threats> Accessed: 2023-10-25.
- [92] Brij B Gupta and Omkar P Badve. 2017. Taxonomy of dos and ddos attacks and desirable defense mechanism in a cloud computing environment. *Neural Computing and Applications* 28, 12 (2017), 3655–3682.
- [93] Saram Han and Chris K Anderson. 2020. Customer motivation and response bias in online reviews. *Cornell Hospitality Quarterly* 61, 2 (2020), 142–153.
- [94] Julie M Haney and Susanne M Furman. 2020. Work in progress: Towards usable updates for smart home devices. In *International Workshop on Socio-Technical Aspects in Security and Trust*. Springer, 107–117.
- [95] Julie M Haney and Susanne M Furman. 2023. Smart home device loss of support: Consumer perspectives and preferences. In *International Conference on Human-Computer Interaction*. Springer, 492–510.
- [96] Julie M Haney, Susanne M Furman, and Yasemin Acar. 2020. Smart home security and privacy mitigations: Consumer perceptions, practices, and challenges. In *International Conference on Human-Computer Interaction*. Springer, 393–411.
- [97] Sherry He, Brett Hollenbeck, and Davide Proserpio. 2022. The market for fake reviews. *Marketing Science* (2022).
- [98] Malcolm Heath. 2023. 2023 DDoS Attack Trends. <https://www.f5.com/labs/articles/threat-intelligence/2023-ddos-attack-trends> [Accessed: October 10, 2025].
- [99] Cristian Hesselman and Ramin Yazdani. 2020. DDoS Clearing House for Europe Cross-sector Pilot Demo. <https://www.sidnlabs.nl/downloads/2deJudioEsd0oFWufTXdV9/099fa8c92f7d601e0669bec73b2fa272/NEW-20200123-CONCORDIA-T3.2-demo-review-final.pdf> [Accessed: October 10, 2025].
- [100] David W. Hosmer, Trina A. Hosmer, Saskia le Cessie, and Stanley Lemeshow. 1997. A comparison of goodness-of-fit tests for the logistic regression model. *Statistics in medicine* 16 9 (1997), 965–80.
- [101] H.R.1668. 2019-2020. IoT Cybersecurity Improvement Act of 2020. <https://www.congress.gov/bill/116th-congress/house-bill/1668> Accessed: 2024-11-14.
- [102] Nan Hu, Paul A Pavlou, and Jie Jennifer Zhang. 2017. On self-selection biases in online product reviews. *MIS Q.* 41, 2 (2017), 449–471.
- [103] Md. Hussain and Ishtiaq Mahmud. 2019. Pymannkendall: A python package for non parametric mann kendall family of trend tests. *Journal of Open Source Software* 4, 39 (25 7 2019), 1556. <https://doi.org/10.21105/joss.01556>

- [104] GreyNoise Intelligence. 2024. Greynoise intelligence discovers zero-day vulnerabilities in live streaming cameras with the help of ai. <https://www.greynoise.io/blog/greynoise-intelligence-discovers-zero-day-vulnerabilities-in-live-streaming-cameras-with-the-help-of-ai> Accessed: February 25, 2025.
- [105] Privacy International. 2022. Privacy international research shows that smart device security updates fail to meet consumers' Expectations. <https://privacyinternational.org/press-release/4964/privacy-international-research-shows-smart-device-security-updates-fail-meet> Accessed: February 24, 2025.
- [106] Privacy International. 2022. We looked into the software support practices for 5 of the most popular smart devices (and the results may disappoint you). <https://privacyinternational.org/report/4965/we-looked-software-support-practices-5-most-popular-smart-devices-and-results-may> Accessed: February 24, 2025.
- [107] IoT Analytics. 2024. *Insights release: Number of connected IoT devices (2024)*. Technical Report. <https://iot-analytics.com/wp-content/uploads/2024/09/INSIGHTS-RELEASE-Number-of-connected-IoT-devices-vf.pdf> PDF insights release providing updated figures and trends on global IoT device connectivity.
- [108] ITSecurityGuru. 2018. Proactive vs. Reactive: Which is better for ddos defence? <https://www.itsecurityguru.org/2018/01/29/proactive-vs-reactive-better-ddos-defence/> [Accessed: October 10, 2025].
- [109] James A. Jerkins. 2017. Motivating a market or regulatory solution to IoT insecurity with the Mirai botnet code. In *2017 IEEE 7th Annual Computing and Communication Workshop and Conference, CCWC 2017*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/CCWC.2017.7868464>
- [110] Ron De Jesus. 2020. How to operationalize privacy by design. <https://iapp.org/news/a/how-to-operationalize-privacy-by-design> Accessed: 2024-11-14.
- [111] 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.
- [112] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of targets under attack: A macroscopic characterization of the DoS ecosystem. (2017), 100–113.
- [113] Mattijs Jonker, Anna Sperotto, Roland van Rijswijk-Deij, Ramin Sadre, and Aiko Pras. 2016. Measuring the adoption of ddos protection services. In *Proceedings of the 2016 Internet Measurement Conference*. 279–285.

- [114] Davino Mauro Junior, Luis Melo, Hao Lu, Marcelo d'Amorim, and Atul Prakash. 2019. A study of vulnerability analysis of popular smart devices through their companion apps. In *2019 IEEE Security and Privacy Workshops (SPW)*. IEEE, 181–186.
- [115] Georgios Kambourakis, Tassos Moschos, Dimitris Geneiatakis, and Stefanos Gritzalis. 2007. Detecting dns amplification attacks. In *International workshop on critical information infrastructures security*. Springer, 185–196.
- [116] Mohammad Karami and Damon McCoy. 2013. Understanding the emerging threat of ddos-as-a-service. In *6th {USENIX} Workshop on Large-Scale Exploits and Emergent Threats ({LEET} 13)*.
- [117] Mohammad Karami, Youngsam Park, and Damon McCoy. 2016. Stress testing the booters: Understanding and undermining the business of ddos services. In *Proceedings of the 25th International Conference on World Wide Web*. 1033–1043.
- [118] Nickson M. Karie, Nor Masri Sahri, Nor Masri Sahri, Wencheng Yang, Craig Valli, Craig Valli, and Victor R. Kebande. 2021. A review of security standards and frameworks for IoT-based smart environments. *IEEE Access* (2021). <https://doi.org/10.1109/access.2021.3109886>
- [119] Erin Kenneally and David Dittrich. 2012. The Menlo report: Ethical principles guiding information and communication technology research. *Available at SSRN 2445102* (2012).
- [120] David R. Anderson Kenneth P. Burnham. 2002. *Model selection and inference: A practical information-theoretic approach*. JSTOR.
- [121] Sye Loong Keoh, Sandeep S Kumar, and Hannes Tschofenig. 2014. Securing the Internet of things: A standardization perspective. *IEEE Internet of things Journal* 1, 3 (2014), 265–275.
- [122] Wolfgang Kerber. 2023. Governance of IoT data: Why the eu data act will not fulfill its objectives. *GRUR International* 72, 2 (2023), 120–135.
- [123] Jan-Peter Kleinhans. 2018. Improving IoT security in the eu. <https://www.stiftung-nv.de/en/publication/improving-iot-security-eu>
- [124] Constantinos Koliass, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. 2017. DDoS in the IoT: Mirai and other botnets. *Computer* 50, 7 (2017), 80–84. <https://doi.org/10.1109/MC.2017.201>
- [125] Daniel Kopp, C. Dietzel, and O. Hohlfeld. 2021. Ddos never dies? An ixp perspective on ddos amplification attacks. *PAM* (2021). [https://doi.org/10.1007/978-3-030-72582-2\\_17](https://doi.org/10.1007/978-3-030-72582-2_17)
- [126] 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.

- [127] Sara Kraemer and Pascale Carayon. 2007. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied ergonomics* 38, 2 (2007), 143–154.
- [128] Brian Krebs. 2025. Krebsonsecurity hit with near-record 6.3 tbps ddos. <https://krebsonsecurity.com/2025/05/krebsonsecurity-hit-with-near-record-6-3-tbps-ddos/> Accessed: 2025-05-26.
- [129] Julie Kristin, Marlina Widiyanti, Muchsin Saggaff Shihab, and Aslamia Rosa. 2024. The influence of online consumer reviews, prices, and consumer trust on online purchase decisions for skincare products on shopee. *Enrichment: Journal of Multidisciplinary Research and Development* 2, 5 (2024).
- [130] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. Amppot: Monitoring and defending against amplification ddos attacks. *RAID* (2015). [https://doi.org/10.1007/978-3-319-26362-5\\_28](https://doi.org/10.1007/978-3-319-26362-5_28)
- [131] Marc Kührer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Exit from Hell? Reducing the impact of amplification DDoS attacks. In *23rd USENIX security symposium (USENIX security 14)*. 111–125.
- [132] 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)*. 1169–1185.
- [133] Lorenz Kustosch, Carlos Gañán, Mattis Van’t Schip, Michel Van Eeten, and Simon Parkin. 2023. Measuring up to (reasonable) consumer expectations: Providing an empirical basis for holding {IoT} manufacturers legally responsible. In *32nd USENIX Security Symposium (USENIX Security 23)*. 1487–1504.
- [134] Sheng-Hsien Lee. 2009. How do online reviews affect purchasing intention? *African Journal of Business Management* 3, 10 (2009), 576–581.
- [135] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research* 18, 17 (2017), 1–5. <http://jmlr.org/papers/v18/16-365.html>
- [136] Eireann Leverett, Richard Clayton, and Ross Anderson. 2017. Standardisation and certification of the ‘Internet of Things’. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, Vol. 2017.
- [137] Fraser Lewis, Adam Butler, and Lucy Gilbert. 2011. A unified approach to model selection using the likelihood ratio test. *Methods in ecology and evolution* 2, 2 (2011), 155–162.

- [138] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. 2019. Keepers of the machines: Examining how system administrators manage software updates for multiple machines. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 273–288.
- [139] Chin-Lung Lin, Sheng-Hsien Lee, and Der-Juinn Horng. 2011. The effects of online reviews on purchasing intention: The moderating role of need for cognition. *Social Behavior and Personality: an international journal* 39, 1 (2011), 71–81.
- [140] Mikhail Lisovich and Stephen Wicker. 2008. Privacy concerns in upcoming residential and commercial demand-response systems. *IEEE Proceedings on Power Systems* 1, 1 (2008), 1–10.
- [141] Yong Liu. 2006. Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue. *Journal of marketing* 70, 3 (2006), 74–89.
- [142] Ya Liu and Hui Wang. 2018. Tracking mirai variants. *Virus Bulletin* (2018), 1–18.
- [143] Franco Loi, Arunan Sivanathan, Hassan Habibi Gharakheili, Adam Radford, and Vijay Sivaraman. 2017. Systematically Evaluating Security and Privacy for Consumer IoT Devices. *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy* (2017). <https://api.semanticscholar.org/CorpusID:607695>
- [144] Qasim Lone, Alisa Frik, Matthew Luckie, Maciej Korczyński, Michel van Eeten, and Carlos Ganán. 2022. Deployment of source address validation by network operators: A randomized control trial. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2361–2378.
- [145] Daniel Lüdecke. 2023. *SjPlot: Data Visualization for Statistics in Social Science*. <https://CRAN.R-project.org/package=sjPlot> R package version 2.8.15.
- [146] Lydia Manikonda, Aditya Deotale, and Subbarao Kambhampati. 2018. What's up with privacy? User preferences and privacy concerns in intelligent personal assistants. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 229–235.
- [147] Sonia San Martín and Carmen Camarero. 2005. Consumer reactions to firm signals in asymmetric relationships. *Journal of Service Research* 8, 1 (2005), 79–97.
- [148] 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.
- [149] Dave McMillen. 2021. Internet of Threats: IoT Botnets drive surge in network attacks. <https://securityintelligence.com/posts/internet-of-threats-iot-botnets-network-attacks/> [Accessed: October 10, 2025].

- [150] Katerina Megas, Barbara Cuthill, and Sarbari Gupta. 2021. *Establishing confidence in IoT device security: How do we get there?(draft)*. Technical Report. National Institute of Standards and Technology.
- [151] 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.
- [152] Markus Miettinen, Samuel Marchal, Ibbad Hafeez, N Asokan, Ahmad-Reza Sadeghi, and Sasu Tarkoma. 2017. IoT sentinel: Automated device-type identification for security enforcement in IoT. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2177–2184.
- [153] Wendy W Moe and Michael Trusov. 2011. The Value of Social Dynamics in Online Product Ratings Forums. *Journal of Marketing Research* 48, 3 (2011), 444–456.
- [154] Tyler Moore, Richard Clayton, and Ross Anderson. 2009. The economics of online crime. *Journal of Economic Perspectives* 23, 3 (2009), 3–20.
- [155] Philipp Morgner and Zinaida Benenson. 2018. Exploring security economics in IoT standardization efforts. *Workshop on Decentralized IoT Security and Standards (DISS) 2018* (2018).
- [156] Philipp Morgner, Felix Freiling, and Zinaida Benenson. 2018. Opinion: Security lifetime labels-overcoming information asymmetry in security of IoT consumer products. In *Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks*. 208–211.
- [157] 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.
- [158] Ken Munro. 2018. Why is consumer IoT insecure? <https://www.pentestpartners.com/security-blog/why-is-consumer-iot-insecure/>. Accessed: 2025-09-28.
- [159] Ken Munro. 2018, note = [Accessed: October 10, 2025]. Why is consumer IoT insecure? <https://www.pentestpartners.com/security-blog/why-is-consumer-iot-insecure/>
- [160] Roberto Musotto and David S. Wall. 2020. More Amazon than Mafia: Analysing a DDoS stresser service as organised cybercrime. *Trends in Organized Crime* (2020). <https://doi.org/10.1007/s12117-020-09397-5>
- [161] Yana Myroshnyk. 2024. *State of IoT summer 2024*. Technical Report. IoT Analytics GmbH. <https://iot-analytics.com/product/state-of-iot-summer-2024/> 171-page report providing a comprehensive assessment of the IoT market, including forecasts, trends, and analysis based on Q1 and Q2 2024 data.

- [162] 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.
- [163] Majid Nasirinejad and Srinivas Sampalli. 2023. Evaluating consumer behavior, decision making, risks, and challenges for buying an IoT product. *IoT* 4, 2 (2023), 78–94.
- [164] Netherlands Authority for Consumers and Markets (ACM). 2023. Amazon.NL now informs customers about updates when purchasing smart devices. <https://www.acm.nl/en/publications/amazonnl-now-informs-customers-about-updates-when-purchasing-smart-devices> Accessed: 2025-04-15.
- [165] NetScout. 2022. Botnets Multiply and Level Up. <https://www.netscout.com/threatreport/botnets-multiply-and-level-up/> [Accessed: October 10, 2025].
- [166] Jakob Nielsen. 1994. *Usability engineering*. Morgan Kaufmann.
- [167] NIST. 2022. NIST Issues Guidance on Software, IoT Security and Labeling. <https://www.nist.gov/news-events/news/2022/02/nist-issues-guidance-software-iot-security-and-labeling> Accessed: 2024-11-14.
- [168] Nokia. 2022. Nokia deepfield network intelligence report ddos in 2021. [https://onestore.nokia.com/asset/211059?\\_ga=2.140826161.227459188.1657444403-1091153497.1656679580](https://onestore.nokia.com/asset/211059?_ga=2.140826161.227459188.1657444403-1091153497.1656679580) [Accessed: October 10, 2025].
- [169] A. Noroozian, M. Korczynski, C. Ganan, D. Makita, K. Yoshioka, and M. Van Eeten. 2016. Who Gets the Boot? Analyzing Victimization by DDos-as-a-Service. *RAID* (2016). [https://doi.org/10.1007/978-3-319-45719-2\\_17](https://doi.org/10.1007/978-3-319-45719-2_17)
- [170] Arman Noroozian, Maciej Korczyński, Carlos Hernandez Gañan, Daisuke Makita, Katsunari Yoshioka, and Michel Van Eeten. 2016. Who gets the boot? Analyzing victimization by ddos-as-a-service. In *Research in Attacks, Intrusions, and Defenses: 19th International Symposium, RAID 2016, Paris, France, September 19-21, 2016, Proceedings 19*. Springer, 368–389.
- [171] Arman Noroozian, Elsa Turcios Rodriguez, Elmer Lastdrager, Takahiro Kasama, Michel Van Eeten, and Carlos H Gañán. 2021. Can ISPs help mitigate IoT malware? A longitudinal study of broadband ISP security efforts. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 337–352.
- [172] Norbert Nthala and Ivan Flechais. 2018. Informal support networks: An investigation into home data security practices. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018*. 63–82.

- [173] Robert M O'Brien. 2017. Dropping highly collinear variables from a model: Why it typically is not a good idea. *Social Science Quarterly* 98, 1 (2017), 360–375.
- [174] Krebs on Security. 2018. Naming & shaming web polluters: Xiongmai. <https://krebsonsecurity.com/2018/10/naming-shaming-web-polluters-xiongmai/> Accessed: 2024-11-14.
- [175] Charlie Osborne. 2020. NCA launches UK ad campaign to divert kids searching for cybercrime tools. <https://www.zdnet.com/article/nca-launches-ad-campaign-to-divert-kids-searching-for-cybercrime-tools/> [Accessed: October 10, 2025].
- [176] Do-Hyung Park, Jumin Lee, and Ingoo Han. 2007. The effect of online consumer reviews on consumer purchasing intention: The moderating role of involvement. *International journal of electronic commerce* 11, 4 (2007), 125–148.
- [177] 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.
- [178] Roberto Perdisci, Thomas Papastergiou, Omar Alrawi, and Manos Antonakakis. 2020. Iotfinder: Efficient large-scale identification of IoT devices via passive dns traffic analysis. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 474–489.
- [179] Pleuni. 2024. Top 10 online stores in the Netherlands. *Ecommerce News* (June 2024). <https://ecommercenews.eu/top-10-online-stores-in-the-netherlands/>
- [180] 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.
- [181] Vijay Prakash, Sicheng Xie, and Danny Yuxing Huang. 2022. Inferring Software Update Practices on Smart Home IoT Devices Through User Agent Analysis. In *Proceedings of the 2022 ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses*. 93–103.
- [182] Alvaro Puig. 2024. How long will your smart device get software updates? It's hard to know. <https://consumer.ftc.gov/consumer-alerts/2024/11/how-long-will-your-smart-device-get-software-updates-its-hard-know> Accessed: 2025-03-31.
- [183] S. Rivera Pérez, M. van Eeten, and C. H. Gañán. 2024. Patchy performance? Uncovering the vulnerability management practices of IoT-centric vendors. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 157–157. <https://doi.org/10.1109/SP54263.2024.00154>

- [184] R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/> ISBN 3-900051-07-0.
- [185] Emilee Rader and Rick Wash. 2015. Identifying patterns in informal sources of security information. *Journal of Cybersecurity* 1, 1 (2015), 121–144.
- [186] 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)*. 1–17.
- [187] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2016. How I learned to be secure: A census-representative survey of security advice sources and behavior. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 666–677.
- [188] Santiago Ruano Rincón, Sandrine Vaton, Antoine Beugnard, and Serge Garlatti. 2015. Semantics based analysis of botnet activity from heterogeneous data sources. In *2015 International Wireless Communications and Mobile Computing Conference (IWCMC)*. IEEE, 391–396.
- [189] Elsa Rodríguez, Arman Noroozian, Michel van Eeten, and Carlos Gañán. 2021. Superspreaders: Quantifying the role of IoT manufacturers in device infections. In *20th Workshop on the Economics of Information Security (WEIS)*.
- [190] Everett M Rogers, Arvind Singhal, and Margaret M Quinlan. 2014. Diffusion of innovations. In *An integrated approach to communication theory and research*. Routledge, 432–448.
- [191] Alvin E Roth. 2015. *Who gets what—And why: The new economics of matchmaking and market design*. Houghton Mifflin Harcourt.
- [192] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Hamed Haddadi, Daniel J Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. 2020. A haystack full of needles: Scalable detection of IoT devices in the wild. In *Proceedings of the ACM Internet Measurement Conference*. 87–100.
- [193] Isabella Sansone. 2021. The Damaging Impacts of DDoS Attacks. <https://www.corero.com/the-damaging-impacts-of-ddos-attacks/>
- [194] Matias RP Santos, Rossana MC Andrade, Danielo G Gomes, and Arthur C Callado. 2018. An efficient approach for device identification and traffic classification in IoT ecosystems. In *2018 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 00304–00309.
- [195] SB-327. 2017-2018. Information privacy: Connected devices. [https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill\\_id=201720180SB327](https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201720180SB327) Accessed: 2024-11-14.

- [196] Bruce Schneier. 2018. *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company.
- [197] Bruce Schneier. 2023. We need to save the internet from the internet of things. <https://www.vice.com/en/article/ezpq3m/we-need-to-save-the-internet-from-the-internet-of-things> Accessed on November 29, 2023.
- [198] Martin Schulz and Lloyd A Jobe. 2001. Codification and tacitness as knowledge management strategies: An empirical exploration. *The Journal of High Technology Management Research* 12, 1 (2001), 139–165.
- [199] KungHsin Shao. 2012. The effects of controversial reviews on product sales performance: The mediating role of the volume of word of mouth. *International Journal of Marketing Studies* 4, 4 (2012), 32.
- [200] Rahul Anand Sharma, Elahe Soltanaghaei, Anthony Rowe, and Vyas Sekar. 2022. Lumos: Identifying and localizing diverse hidden {IoT} devices in an unfamiliar environment. In *31st USENIX Security Symposium (USENIX Security 22)*. 1095–1112.
- [201] Yun Shen and Pierre-Antoine Vervier. 2019. IoT security and privacy labels. In *Privacy Technologies and Policy: 7th Annual Privacy Forum, APF 2019, Rome, Italy, June 13–14, 2019, Proceedings 7*. Springer, 136–147.
- [202] Shopify. 2024. Global ecommerce sales: What the latest data shows. <https://www.shopify.com/blog/global-ecommerce-sales> Accessed: 2025-05-01.
- [203] Omer Shwartz, Yael Mathov, Michael Bohadana, Yuval Elovici, and Yossi Oren. 2018. Reverse engineering IoT devices: Effective techniques and methods. *IEEE Internet of Things Journal* 5, 6 (2018), 4965–4976.
- [204] Michael Spence. 1973. Job market signaling. *The Quarterly Journal of Economics* 87, 3 (1973), 355–374. <http://www.jstor.org/stable/1882010>
- [205] Statista. 2023. Consumer IoT - Worldwide. <https://www.statista.com/outlook/tmo/internet-of-things/consumer-iot/worldwide>
- [206] Curtis Steward Jr, Luay A Wahsheh, Aftab Ahmad, Jonathan M Graham, Cheryl V Hinds, Aurelia T Williams, and Sandra J DeLoatch. 2012. Software security: The dangerous afterthought. In *2012 Ninth International Conference on Information Technology-New Generations*. IEEE, 815–818.
- [207] HP Online Store. 2023. The state of printer security. <https://www.hp.com/th-en/shop/tech-takes/post/the-state-of-printer-security> Accessed: 2023-10-28.
- [208] Samaneh Tajalizadehkhoob, Maciej Korczyński, Arman Noroozian, Carlos Ganán, and Michel Van Eeten. 2016. Apples, oranges and hosting providers: Heterogeneity and security in the hosting market. In *NOMS 2016-2016 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 289–297.

- [209] Samaneh Tajalizadehkhoob, Tom Van Goethem, Maciej Korczyński, Arman Noroozian, Rainer Böhme, Tyler Moore, Wouter Joosen, and Michel van Eeten. 2017. Herding vulnerable cats: A statistical approach to disentangle joint responsibility for web security in shared hosting. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 553–567.
- [210] Rui Tanabe, Tatsuya Tamai, Akira Fujita, Ryoichi Isawa, Katsunari Yoshioka, Tsutomu Matsumoto, Carlos Gañán, and Michel van Eeten. 2020. Disposable botnets: Examining the anatomy of IoT botnet infrastructure. In *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES '20)*. Association for Computing Machinery, New York, NY, USA, Article 7, 10 pages.
- [211] Leonie Maria Tanczer, Isabel López-Neira, and Simon Parkin. 2021. ‘I feel like we’re really behind the game’: Perspectives of the United Kingdom’s intimate partner violence support sector on the rise of technology-facilitated abuse. *Journal of Gender-Based Violence* (2021).
- [212] Ars Technica. 2020. Cache issue causes xiaomi cameras to show other people’s camera feeds. <https://arstechnica.com/gadgets/2020/01/cache-issue-causes-xiaomi-cameras-to-show-other-peoples-camera-feeds/> Accessed: 2023-10-25.
- [213] Kurt Thomas, Danny Huang, David Wang, Elie Bursztein, Chris Grier, Thomas J Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. 2015. Framing dependencies introduced by underground commoditization. *Workshop on the Economics of Information Security (WEIS)* (2015).
- [214] Alethea Toh. 2021. Azure ddos protection—2020 year in review. <https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2020-year-in-review/> [Accessed: October 10, 2025].
- [215] Alexandra Topping. 2019. Ring hackers reportedly watching and talking to strangers in home cameras. *The Guardian* (13 December 2019). <https://www.theguardian.com/technology/2019/dec/13/ring-hackers-reportedly-watching-talking-strangers-in-home-cameras> Accessed: 2025-05-01.
- [216] Manuel Trenz and Benedikt Berger. 2013. Analyzing online customer reviews - An interdisciplinary literature review and research agenda. (2013).
- [217] 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>. [Accessed: October 10, 2025].
- [218] UK Department for Digital, Culture, Media & Sport (DCMS). 2021. Regulating consumer smart product cyber security - government response. <https://www.gov.uk/government/publications/regulating-consumer->

- [smart-product-cyber-security-government-response](#). [Accessed: October 10, 2025].
- [219] Karen Van Dam, Shaul Oreg, and Birgit Schyns. 2008. Daily work contexts and resistance to organisational change: The role of leader-member exchange, development climate, and change process characteristics. *Applied psychology* 57, 2 (2008), 313–334.
- [220] Dirk van der Linden, Matthew Edwards, Irit Hadar, and Anna Zamansky. 2020. Pets without PETs: On pet owners’ under-estimation of privacy concerns in pet wearables. *Proc. Priv. Enhancing Technol.* 2020, 1 (2020), 143–164.
- [221] Veerle van Harten, Carlos Hernández Gañán, Michel van Eeten, and Simon Parkin. 2023. Easier said than done: The failure of top-level cybersecurity advice for consumer IoT devices. arXiv:cs.CR/2310.00942 <https://arxiv.org/abs/2310.00942>
- [222] Tanvi Vats, Neelima Sailaja, and Fabiana Anselmo Polido Lopes. 2024. Exploration of user perspectives around software and data-related challenges associated with IoT repair and maintenance against obsolescence: User study on software and data interactions and considerations for IoT repair and maintenance against obsolescence. In *Proceedings of the 13th Nordic Conference on Human-Computer Interaction*. 1–17.
- [223] Swaathi Vetrivel, Brennen Bouwmeester, Michel van Eeten, and Carlos H Gañán. 2024. IoT market dynamics: An analysis of device sales, security and privacy signals, and their interactions. In *33rd USENIX Security Symposium (USENIX Security 24)*. 7031–7048.
- [224] Swaathi Vetrivel, Brennen Bouwmeester, Michel van Eeten, and Carlos H Gañán. 2024. IoT market dynamics: An analysis of device sales, security and privacy signals, and their interactions. In *33rd USENIX Security Symposium (USENIX Security 24)*. 7031–7048.
- [225] Swaathi Vetrivel, Veerle Van Harten, Carlos H Gañán, Michel Van Eeten, and Simon Parkin. 2023. Examining consumer reviews to understand security and privacy issues in the market of smart home devices. In *32nd USENIX Security Symposium (USENIX Security 23)*. 1523–1540.
- [226] Daniel Wagner, Daniel Kopp, M. Wichtlhuber, C. Dietzel, O. Hohlfeld, Georgios Smaragdakis, and A. Feldmann. 2021. United we stand: Collaborative detection and mitigation of amplification ddos attacks at scale. *CCS* (2021). <https://doi.org/10.1145/3460120.3485385>
- [227] Xiaopan Wang, Junpeng Guo, Yi Wu, and Na Liu. 2020. Emotion as signal of product quality: Its effect on purchase decision based on online customer reviews. *Internet Research* 30, 2 (2020), 463–485.

- [228] David Warburton. 2021. Ddos attack trends for 2020. <https://www.f5.com/labs/articles/threat-intelligence/ddos-attack-trends-for-2020> [Accessed: October 10, 2025].
- [229] Meredydd Williams, Jason RC Nurse, and Sadie Creese. 2017. Privacy is the boring bit: User perceptions and behaviour in the Internet-of-Things. In *2017 15th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 181–18109.
- [230] Worldometer. 2025. World population clock: Live. <https://www.worldometers.info/world-population/> Accessed: 2025-05-01.
- [231] Shundan Xiao, Jim Witschey, and Emerson Murphy-Hill. 2014. Social influences on secure development tool adoption: Why security tools spread. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 1095–1106.
- [232] Majid Yar. 2005. The Novelty of ‘Cybercrime’ An Assessment in Light of Routine Activity Theory. *European Journal of Criminology* 2, 4 (2005), 407–427.
- [233] Narges Yousefnezhad, Avleen Malhi, and Kary Främling. 2020. Security in product lifecycle of IoT devices: A survey. *Journal of Network and Computer Applications* 171 (2020), 102779.
- [234] Lingjing Yu, Bo Luo, Jun Ma, Zhaoyu Zhou, and Qingyun Liu. 2020. You are what you broadcast: Identification of mobile and IoT devices from (public) wifi.. In *USENIX Security Symposium*. 55–72.
- [235] Eric Zeng, Shrirang 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.
- [236] Jiaming Zhan, Han Tong Loh, and Ying Liu. 2009. Gather customer concerns from online product reviews—a text summarization approach. *Expert Systems with Applications* 36, 2 (2009), 2107–2115.
- [237] Yueyue Zhang, Cheng Zhang, and Yunjie Xu. 2021. Effect of data privacy and security investment on the value of big data firms. *Decision Support Systems* 146 (2021), 113543.
- [238] Binbin Zhao, Shouling Ji, Wei-Han Lee, Changting Lin, Haiqin Weng, Jingzheng Wu, Pan Zhou, Liming Fang, and Raheem Beyah. 2020. A large-scale empirical study on the vulnerability of deployed IoT devices. *IEEE Transactions on Dependable and Secure Computing* 19, 3 (2020), 1826–1840.
- [239] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User perceptions of smart home IoT privacy. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–20.

- [240] Feng Zhu and Xiaoquan Zhang. 2010. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of marketing* 74, 2 (2010), 133–148.
- [241] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. 2021. Asdb: A system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC '21)*. Association for Computing Machinery, New York, NY, USA, 703–719. <https://doi.org/10.1145/3487552.3487853> [Accessed: October 10, 2025].

# A

## APPENDIX

### A.1. SEARCH TERMS USED FOR EACH OF THE FOUR DEVICE TYPES

**Surveillance Systems:** Surveillance camera, Network Camera, IP Camera, Security Camera, Dome Camera and DVR/NVR

**Smart home hub:** Smart home hub, Smart home control panel, Smart home automation system

**Set-top box:** Digital set top box, IP set top box

**Router:** Router, WiFi repeater

### A.2. SEARCH TERMS USED FOR IDENTIFYING S&P RELATED CUSTOMER REVIEWS BY CATEGORY

**Configuration and Authentication:** setup, configure, profile, default, control, 2FA, authentication, password

**Access and Storage:** access, manipulate, watch, track, record, log, data

**Encryption and Security:** encrypt, encryption, ssl, protocol, secure

**Privacy:** privacy, private, personal, trust, open

**Attack:** hack, attack, fraudsters, spy, steal, blackmail, criminals, cutoff

**Patches and updates:** patch, uptodate, update, firmware, vulnerability, risk, safe, protect

### A.3. RESULTS OF LDA

#### LDA for Surveillance Systems

**Topic 1:** night, vision, product, quality, picture, light, day, time, image, work || **Topic 2:** app, phone, time, monitor, quality, video, home, view, picture, work || **Topic 3:** system, quality, nvr, image, security, poe, setup, setting, video, picture || **Topic 4:** network, device, connection, work, setup, router, app, internet, issue, access || **Topic 5:** unit, battery, ring, model, number, year, doorbell, resolution, zone, month || **Topic 6:** cloud, storage, sub-

scription, video, window, stream, service, option, plan, year || **Topic 7:** motion, detection, notification, alert, time, sensitivity, record, video, setting, alarm || **Topic 8:** cam, brand, contact, color, noise, today, stuff, good, audio, condition || **Topic 9:** cable, power, wire, ethernet, wall, box, plug, screw, plastic, plate || **Topic 10:** card, video, sd, record, app, footage, recording, memory, playback, file || **Topic 11:** software, car, pc, door, computer, web, hardware, people, foot, interface || **Topic 12:** support, customer, service, issue, problem, email, help, tech, replacement, update

#### LDA for Routers

**Topic 1:** extender, unit, range, room, instruction, wifi, work, light, install, plug || **Topic 2:** network, setup, system, app, home, mesh, point, access, port, cable || **Topic 3:** connection, performance, laptop, bit, set, video, quality, eero, phone, work || **Topic 4:** speed, signal, house, internet, coverage, floor, strength, wifi, drop, test || **Topic 5:** support, time, work, day, money, service, hour, tech, week, customer || **Topic 6:** issue, price, month, review, year, band, problem, time, model, day

#### LDA for Hubs

**Topic 1:** device, time, app, work, product, tv, control, button, setup, hub || **Topic 2:** music, speaker, sound, play, love, sound\_quality, room, alarm, quality, question

#### LDA for Set-top Boxes

**Topic 1:** tv, box, fire, device, app, stick, work, product, control, issue || **Topic 2:** channel, tv, time, record, program, device, unit, guide, cable, recording

#### LDA for Once-vulnerable Routers

**Topic 1:** network, setup, connection, speed, work, range, feature, access, signal, option || **Topic 2:** issue, time, internet, connection, day, problem, cable, work, unit, support

## A.4. RESULTS OF ANCHORED COREX

**Topic 1** (23.8%) : *anchor words: setup, configure, control* - setup, password, easy setup, setup easy, easy, initial setup, configure, username password, username, camera setup || **Topic 2** (19.8%) : (*anchor words: access, watch, record*) - record, camera, motion, night, detection, quality, vision, night vision, motion detection, video || **Topic 3** (11.1%) : (*anchor words: encrypt, secure, protocol*) - secure, protocol, encrypt, address, iris, blue iris, onvif || **Topic 4** (15.6%) : (*anchor words: open, trust, personal*) - open, trust, personal, open source, time open, open camera, personal data, camera open, open door, seal || **Topic 5** (11.5%) : (*anchor words: steal, hack, spy*) - hack, steal, price, easily, attach, small, hole, recommend, outside, white || **Topic 6** (8.9%) : (*anchor words: protect, update, firmware*) - update, firmware, firmware update, update firmware, latest, latest firmware, version, update review, upgrade, firmware upgrade || **Topic 7** (6.5%) : *no anchors* - support, work, time, issue, review, try, problem, email, reset, contact || **Topic 8** (2.8%) : *no anchors* - network, connect, setting, devices, cable, power, connection, feature, point, plug

## A.5. CODEBOOK FROM INDUCTIVE CODING

**Firmware updates:** fw-security-patch, fw-auto-update, exclusive-fw-update, fw-available, fw-buggy, fw-comparision, fw-language, fw-latest, fw-update-caused-issue, fw-update-didn't-fix-issue, fw-update-difficult, fw-update-for-feature, fw-update-frequent-good, fw-update-had-feature, fw-update-kludgy, fw-update-might-cause-issue, fw-update-might-

fix, fw-update-not-available, fw-update-timing, fw-update-to-solve-issue, fw-update-took-longer, fw-update-when-setting-up, fw-updates-back-to-back-annoying

**Security during setup:** setup-password, setup-qrcode-easy, setup-qrcode-strange, setup-ssid-not-hidden, setup-too-simple, setup-wps, setup-wps-easy

**Factual Information and advice:** chinese-servers-distrust, contacts-mfg-server, fw-update-not-available-chinese, modem-too-secure, remote-access-cust-support, rtsp-pw-plaintext, security-advice, security-comment, security-protocol, unconcerned-about-vulnerabilities, used-device-pwd-already-set, vulnerable-to-chinese-hackers, woods-tablet

**Negative opinion on security features:** 2fa-missing, guest-network-doesn't-isolate, has-security-vulnerability, I-was-hacked-stories, limited-security, remote-access-uncomfortable, security-concern, unhappy-about-security, unsecure-device-alert

**Usability of security features:** 2FA-useless, can't-change-password, can't-setup-security, conflict-with-wpa3-wps-settings, encrypted-affects-playback, encryption-program-marketing, extra-security-on-trial, locked-out-cos-password, no-password-good, password-not-available, password-reset-thru-customer-support, passwords-cumbersome, security-options-at-launch, share-admin-privileges, share-password-with-qr, unhappy-about-update-for-privacy

**Positive opinion of security features:** 2fa-better-security, configured-for-security-so-not-worried, ddos-protection-good, firewall-configuration-not-needed, firewall-good, good-securitywise, guest-network-for-security, malicious-activity-monitoring-valuable, notification-when-new-device-joins, secure-video-sharing

**Data capitalism and privacy:** app-permissions, cloud-not-needed, face-recognition-good-for-privacy, not-always-listening-is-plus, privacy-concern, privacy-mode-in-camera



# B

## APPENDIX

### B.1. LIST OF SUB-TESTS FOR EACH DEVICE TYPE

**Table B.1:** List of test categories for each device type

IP Cameras	Smart Watches	Smart Printers	Smart Speakers
Password Policy	Standard Settings	Setup	Data Security
Standard Installation	Encryption	Access Controls	Decommissioning
Android App	Factory Reset	Password Policy	Password Policy
iOS App	Man-in-the-Middle	Updates	Network Security
Updates	Updates	Permissions	Update Policy
Known Vulnerabilities	Password Policy	Encryption	iOS App
	Account Management	Authentication	Android App
	Data Minimization	Known Vulnerabilities	Privacy Policy
	Options	Decommissioning	
	Choice Consequences		
	Vulnerability Hotline		

## B.2. PRICE COMPARISON BETWEEN CB TESTED AND NON-TESTED IOT DEVICES

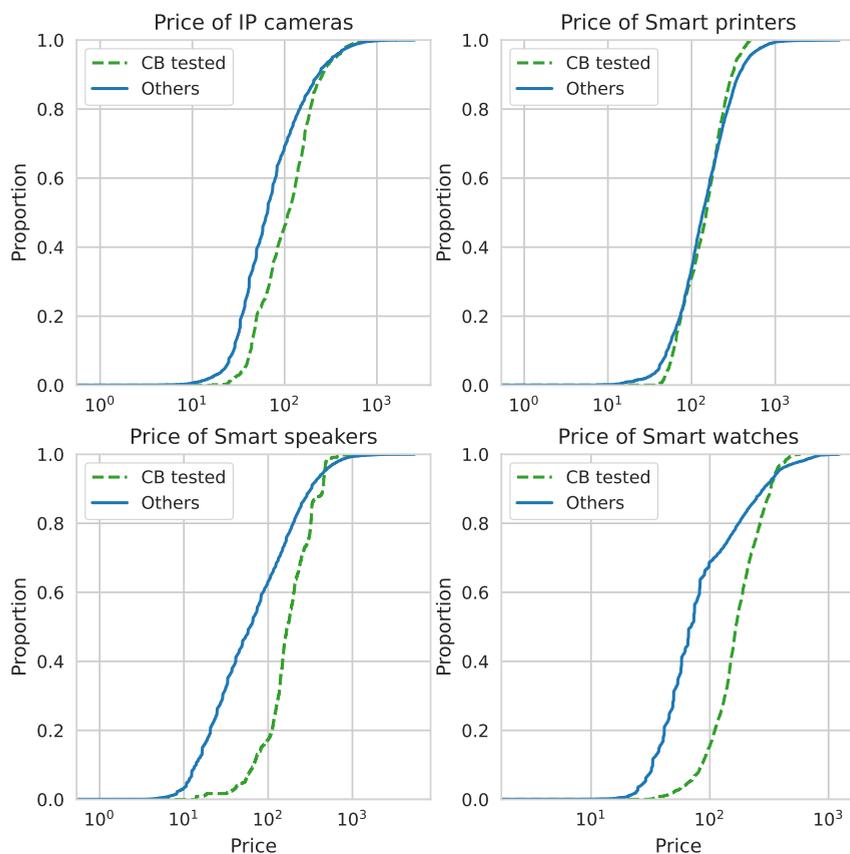


Figure B.1: CDF of price of devices tested by the Consumentenbond and those not

## B.3. GLM VARIABLE SELECTION

Generalised Variance Inflation Factors

	GVIF	Df	GVIF*(1/(2*Df))
Device type	4.623	3	1.291
Average price	1.534	1	1.239
Expert S&P Rating	2.881	1	1.698
Expert Overall Rating	1.804	1	1.343
Winkel user rating	1.345	1	1.596
Winkel total reviews	1.116	1	1.057
Amazon user rating	1.463	1	1.209
Amazon total reviews	1.071	1	1.034

### B.4. CORRELATION BETWEEN CB DEVICE RATINGS AND AVERAGE RATINGS ON AMAZON AND WINKEL

Device type	Spearman correlation	
	Expert Overall rating vs	
	Amazon user rating	Winkel user rating
IP Cameras	0.144	0.141
Smart printers	0.159	0.054
Smart speakers	0.664***	0.289
Smart watches	0.360***	0.124
Aggregate	0.221***	0.127**

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

### B.5. GLM RESULTS WITH UPDATE STATUSES

Table B.2: Mixed Effect Negative Binomial Model with Update Statuses

Dependent variable: Total sales			
Predictors	Incidence Rate Ratios		p
(Intercept)	1807		< 0.001
Avg Price	0.74		0.001
Expert SnP Rating	1.63		< 0.001
Expert Overall Rating	1.23		0.014
Winkel User Rating	1.42		< 0.001
Amazon User Rating	1.15		0.029
Winkel Total Reviews	2.09		< 0.001
Amazon Total Reviews	1.25		0.006
Bol update status [Partly Valid]	1.1		0.768
Bol update status [Unavailable]	0.79		0.444
Bol update status [Valid]	1.69		0.065
Amazon update status [Unavailable]	0.66		0.042
Amazon update status [Valid]	0.87		0.59
Avg Price × Expert SnP Rating	0.91		0.333
Avg Price × Expert Overall Rating	1.15		0.125
Expert SnP Rating × Overall Rating	1.04		0.702
(Avg Price × Expert SnP Rating) × Expert Overall Rating	0.89		0.184
Random Effects			
$\sigma^2$	0.63		
$\tau_{00}$ DeviceType	0.32		
ICC	0.34		
N DeviceType	4		
Observations	302		
Marginal R2 / Conditional R2	0.588/0.727		



# C

## APPENDIX

## C.1. BETA REGRESSION ESTIMATES

**Table C.1:** Beta Regressions Estimates for all Models

<b>Dependent variable: Proportion of Positive S&amp;P Features</b>							
	<b>Model 0</b>	<b>Model 1</b>	<b>Model 2</b>	<b>Model 3</b>	<b>Model 4</b>	<b>Model 5</b>	<b>Model 6</b>
<i>Predictors</i>	<i>IRR</i>						
(Intercept)	0.76 (< 0.001)	0.89 (< 0.001)	0.89 (< 0.001)	1.12 (0.011)	1.09 (0.057)	1.11 (< 0.097)	1.46 (< 0.001)
DeviceType [Smart Printers]		0.79 (< 0.001)	0.79 (< 0.001)	0.56 (< 0.001)	0.58 (< 0.001)	0.56 (< 0.001)	0.52 (< 0.001)
DeviceType [Smart Speakers]		0.84 (0.073)	0.84 (0.078)	0.87 (0.134)	1.09 (0.444)	1.08 (0.504)	0.9 (0.420)
Price			0.99 (0.667)	1.01 (0.598)	1.00 (0.785)	1.01 (0.730)	1.01 (0.606)
Number of Devices				1.21 (< 0.001)	1.21 (< 0.001)	1.21 (< 0.001)	1.13 (< 0.001)
Manufacturer Size					0.94 (0.002)	0.94 (0.003)	0.97 (0.06)
Years since Founding						1.01 (0.648)	1.20 (< 0.01)
HQ [China]							0.82 (< 0.001)
HQ [Japan]							0.67 (< 0.001)
HQ [NL]							0.61 (0.008)
HQ [Swe- den]							0.36 (< 0.001)
HQ [Tai- wan]							0.78 (0.075)
<b>Observations</b>	428	428	428	428	428	428	428
<b>R2</b>	0.000	0.072	0.072	0.135	0.145	0.146	0.255
<b>AIC</b>	-892	-931.3	-929.5	-972	-979.6	-977.8	-1101.3
<b>BIC</b>	-883.9	-915	-909.2	-947.6	-951.1	-945.3	-1048.6
<b>LogLik</b>	448	469.6	469.7	492	496.8	496.9	563.7

# ACKNOWLEDGEMENTS

All PhDs start with knowing where to go, but rarely knowing how to get there. You explore, you flounder, you fly, you fall, and you chart your own way through it all. And though it is an individual journey, it is not one you take alone. There are those who walk with you, those who light the way, and those who cheer you on. This is my moment to thank them all.

First and foremost, to my wonderful, wonderful promotors, I am fortunate to have had the best supervisory team one could ask for. Thank you for your unwavering support and for guiding me on the journey from being an engineer to becoming a socio-technical researcher.

Michel, throughout these years, I have been amazed and humbled by your ability to stay on top of the many different things the team does, while providing each of us with insights that shape and structure our research—all with such kindness. Thank you also for hosting us for dinners and gatherings; we owe our strong, cohesive team spirit to your openness and inclusiveness.

Carlos, thank you for your sharp, insightful comments, for lightening any situation with your sense of humour, and for your periodic dire warnings about time running out. You are the person we turn to when it's 10 pm, we have a midnight deadline, and no clue about PC conflicts. Thank you for always being there for everything, from dealing with admin tasks to fine-tuning model parameters to giving career advice.

I would also like to thank Simon for his tremendous help with the first paper and for showing me the ropes of academic writing. I am inspired by your ability to keep abreast of the literature and connect different, seemingly unrelated strands of research.

Data is to research what fuel is to fire, and I am very grateful to Yvo, Eric, and Brennen for sharing valuable data that enabled this research. Yvo and Eric, thank you for finding a way to share the data despite the many challenges and complexities. Brennen, thank you for going out of your way and generously providing the dataset that allowed for what is one of my favourite analyses in this research.

To Shyamala ma'am and Dr Storm, two people who taught and inspired me, without whom I would not have been interested in Economics of Cybersecurity to begin with. Shyamala ma'am, your course on Security and Cryptography during my Bachelor's kindled my interest in cybersecurity. I've always been impressed by your uncanny knack for explaining the complexities of cryptographic protocols simply, without simplifying them. Dr Storm, thank you for your inspiring lectures on Economics and for making the subject come alive in the classroom through your enthusiasm. Once we see the world through an economic lens, we cannot unsee it.

To our wonderful TPM cybersecurity group, thank you for sharing the many ups and downs of the PhD together. Elsa, you were one of the first people I discussed doing a PhD with. From the moment I walked into your office as a Master's student to ask about it, to almost five years later, figuring out when the layman's talk of the defence is, your

support has been immeasurable. Thank you for your warmth, friendship, and guidance, and for always reminding us that we are doing better than we think we are. Radu, thank you for always being there to help with anything server-related with six sigma guarantees on time. It was a wonderful learning experience to work with you and Yury. I am also grateful to Yury for showing me what an organised notebook looks like and inspiring me to pay more attention to detail. Armaan, thank you for all your support with the longest research study I have undertaken so far. Rolf, thank you for showing us the tricks of the trade and the importance of sticking to the spirit of a rule, not necessarily the letter. Savvas, presentations without memes will never be the same again. Veerle, my (project) partner—from helping each other with research, to our long conversations about everything under the sun, and relying on your gold-standard organising skills to offset my lapses—this journey has been infinitely more enjoyable with you alongside. Arwa, thank you for being my confidante and for the many heart-to-hearts. Laughing together about the you-know-whats made this journey a lot more fun. Lorenz, from going to our first Usenix together to seeing you handle fatherhood so gracefully, it feels like we've come a long way. Mathew, thanks for being a sport about all the leg-pulling and for holding your own despite it. Your convictions have helped me look at things differently. Szu, my newest office mate and setter of high standards for presentations, I am glad you moved to the warmer side, and I am touched by your earnestness in all that you do. Cécile, thank you for dropping in every once in a while and for the honest conversations that leave me feeling lighter. Ronak, I am deeply inspired by your equanimity in the face of any challenges, big or small. I am really glad you joined the group, and thank you for the Hindi lessons—picture abhi baaki hai meri dost! Natalia, Kate, Xander, Qasim, Sandra, Yana, Aksel, Annebel, Donald, Evi, Fieke, Hugo, Kelvin, and Max, thank you all for being the fabric of this wonderful group and enriching this journey in many ways.

I would also like to thank Joy, Joyce, Jolanda, and others within the secretariat for all the work you do behind the scenes to keep the admin wheels running smoothly.

I am deeply grateful for the unwavering support of my parents and family throughout this journey. Appa, thank you for instilling in me a strong work ethic, dedication, and commitment—not through words, but through your living example. Amma, thank you for being my own personal role model for what feminine strength looks like. Aai and Baba, thank you for being the enthusiastic cheerleaders that you are, celebrating all my achievements, big or small. Ammamma, rest assured, I am finally done studying. Thank you for always being my ray of sunshine on a cloudy day. Shruthi, thank you for being there to hear me out and giving me a healthy dose of your no-nonsense, down-to-earth wisdom laced with just enough 'nakkal' to make it palatable. Adhav and Ria, my not-so-tiny bundles of joy, thank you for lighting up all our lives. To Ganesh Mama, Chithi, Chitappa, Athai, Adi, Tara, Raam, and Rishi, thank you for the encouragement and support, and of course, for ganging up to tease me every chance you get. Appaye and Thatha, thank you for making us feel so special whenever we visited and for always stressing the importance of a good education, 'nalla padichu neriya mark vanganum'. Thatha and Madan Mama, thank you for teaching me to be strong, to dream big, and to keep fighting the good fight. I wish you were around to share this moment with me.

I have also been fortunate to find a family in my friends. Shyam and Titiksha, they say you are an average of the five people you spend the most time with. If that is true, I am

extremely glad and grateful that you two are in the mix. Aarabhi and Lee, thank you for bearing with my long silences and still being in my corner. Krishna and Soundarya, I am so glad you live closer now; our trips have been much-needed getaways to rest, recoup, and recharge. Anya, thank you for always being excited and proud of what I am doing. Now, please move closer so I can bond with Siddhu. Mug, thank you for your unchanging groundedness and for showing me that even when things change, they can stay the same. And for the cutest lil pick-me-ups in Ekam's videos. Luca and Mo, thank you for your support and warmth and for being our 'satsanga'.

And now, to the one who has been by my side every step of the way, believing in me even when my own conviction faltered. Omkar, without you, I would have neither started nor completed this. Thank you. This one is for us!

Finally, and most importantly, I offer my deepest gratitude to my Guru for breathing ease into my life and for giving me the tools to cultivate clarity and balance.

|| Śrī Guru Pādam Samarpayāmi ||

*Swaathi Vetrivel*  
*Leiden, October 2025*



# AUTHORSHIP CONTRIBUTIONS

This dissertation is based on five papers that resulted from collaboration with my co-authors. Their participation and feedback in these studies improved the quality and impact of the research. In the following paragraphs, I outline their individual contributions to each of the studies.

For the first study (see Chapter 2), my co-authors Carlos H. Gañán, Michel van Eeten, and Simon Parkin provided valuable feedback on the study design and analysis. They also helped polish the storyline, improve the draft, and proofread the manuscript. Veerle van Harten and I conducted the thematic analysis to identify the themes articulated in the reviews. Simon Parkin assisted with the qualitative analysis by providing timely feedback and suggestions to improve the process. All the data collection, other analyses, and writing were done by me.

For the second study (Chapter 3), Brennen Bouwmeester provided access to the sales data from Winkel. Carlos H. Gañán and Michel van Eeten helped shape the research idea, provided feedback on the analysis, and improved the quality of the writing. I performed all other data collection, analysis, modeling, and writing.

For the third and fourth studies (Chapters 4 and 5), Carlos H. Gañán and Michel van Eeten helped with the overall research design and scoping, including providing valuable feedback on the storyline and positioning. They also helped refine the draft of the papers. All the data collection, analysis, modeling, and writing were done by me.

For the fifth study (see Chapter 6), Daisuke Makita and Katsunari Yoshioka provided access to the data. Arman Noroozian provided inputs on the analysis and writing. Carlos H. Gañán and Michel van Eeten helped refine the idea, gave feedback on the intermediate results, and helped polish the writing. I conducted the analysis and wrote the first draft.

I am grateful to my co-authors for their support and contributions to these studies. In particular, I would like to thank Simon Parkin for his guidance in the first study and for showing me the ropes of qualitative research. I also owe a huge debt of gratitude to my promotors, Michel and Carlos, for their constant support, insights, and inspiration throughout these studies.



# PUBLICATIONS

- **Vetrivel, S.**, Van Harten, V., Gañán, C. H., Van Eeten, M., & Parkin, S. (2023). Examining consumer reviews to understand security and privacy issues in the market of smart home devices. In 32nd USENIX security symposium (USENIX security 23) (pp. 1523-1540).
- Anghel, R., **Vetrivel, S.**, Rodriguez, E.T., Sameshima, K., Makita, D., Yoshioka, K., Gañán, C. & Zhauniarovich, Y. (2023). Peering into the Darkness: The Use of UTRS in Combating DDoS Attacks. In European Symposium on Research in Computer Security (pp. 23-41). Cham: Springer Nature Switzerland.
- **Vetrivel, S.**, Bouwmeester, B., van Eeten, M., & Gañán, C. H. (2024). IoT Market Dynamics: An Analysis of Device Sales, Security and Privacy Signals, and their Interactions. In 33rd USENIX Security Symposium (USENIX Security 24) (pp. 7031-7048).
- **Vetrivel, S.**, Noroozian, A., Makita, D., Yoshioka, K., van Eeten, M., & Gañán, C. H. (2023). Birds of a Feather? A Comparative Analysis of DDoS Victimization by IoT Botnet and Amplification Attacks. In 22nd Workshop on the Economics of Information Security (WEIS 2023).



## ABOUT THE AUTHOR



Swaathi Vetrivel (1992) was born in Chennai, India. In 2013, she completed her Bachelor's degree in Computer Science and Engineering from the Amrita School of Engineering, India. After graduation, she worked for five years as a Software Engineer at Juniper Networks in Bangalore, developing features across different areas of the router operating system.

In 2018, she moved to the Netherlands to pursue a Master's degree in Management of Technology at TU Delft, with a specialisation in the Economics of Cybersecurity. She started her PhD in 2021 with the Cybersecurity Group within the Faculty of Technology, Policy and Management (TPM) at TU Delft.

Her doctoral research investigated market signals related to IoT security and privacy, identifying mechanisms to encourage more security- and privacy-conscious decisions at the point of purchase. Through empirical studies of e-commerce platforms, consumer behaviour, and manufacturer practices, she examined the information asymmetry in the consumer IoT market and the (lack of) signals available to communicate the security and privacy posture of IoT devices. Her work bridges technical security research with economics, platform governance, and market design, providing a comprehensive view of security and privacy within the IoT market ecosystem.

Swaathi is continuing as a postdoctoral researcher in the TPM Cybersecurity Group at TU Delft. She is currently studying how enterprises can strengthen their cybersecurity strategies by improving fundamentals such as vulnerability management. She also explores how organisations can prepare for the future by identifying AI-driven risks and opportunities. Her research focuses on the risk dynamics and incentive structures that shape organisational security decisions, aiming to reshape enterprise approaches to patching and security maintenance.

In addition, since April 2025, Swaathi is a Research Fellow at The Pranava Institute in India, where she explores pathways towards more sustainable, equitable, and inclusive technology deployment in a rapidly changing digital world. Drawing on her cross-sector experience, she adopts a multidimensional approach to responsible technology deployment—moving beyond simplistic binaries of whether or not to deploy technologies, and instead asking how we can consciously design and deploy them while remaining mindful of their longer-term consequences.



encrypt

update

secure

WEP

WAP

Permission

f r e w

hacked

safe

patch

2FA