

Document Version

Final published version

Citation (APA)

Kustosch, L. F. (2026). *Security by Expectation: Establishing an Empirical Understanding of Reasonable User Expectations in the Internet of Things*. [Dissertation (TU Delft), Delft University of Technology].
<https://doi.org/10.4233/uuid:914073ad-a3d7-45ac-94f3-f64c41cac651>

Important note

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

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

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.

SECURITY BY EXPECTATION

**Establishing an Empirical
Understanding of Reasonable
User Expectations in the Internet of Things**

Lorenz Kustosch

SECURITY BY EXPECTATION:
ESTABLISHING AN EMPIRICAL UNDERSTANDING OF
REASONABLE USER EXPECTATIONS IN THE INTERNET OF THINGS



SECURITY BY EXPECTATION:
ESTABLISHING AN EMPIRICAL UNDERSTANDING OF
REASONABLE USER EXPECTATIONS IN THE INTERNET OF THINGS

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. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Wednesday, 3 June 2026 at 10:00

by

Lorenz Ferdinand KUSTOSCH

This dissertation has been approved by the promotor:

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

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof. dr. M.J.G. van Eeten	Delft University of Technology, promotor
Dr. ir. C. Hernández Gañán	Delft University of Technology, promotor
Dr. S.E. Parkin	Delft University of Technology, copromotor

Independent members:

Prof. dr. ir. P.H.A.J.M. van Gelder	Delft University of Technology
Prof. dr. T.A.P. Metzke	Delft University of Technology
Prof. dr. A. Rashid	University of Bristol
Prof. mr. F.J. Zuiderveen Borgesius	Radboud University
Dr. R.S. van Wegberg	Delft University of Technology

This research was funded by the INTERSECT project (Grant nr. NWA.1160.18.301), financed by the Dutch Research Council (NWO).



Cover by: Ketewan Kustosch

Copyright © 2026 by L.F. Kustosch

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

To Ketí, Mathilda, Karin, and Jürgen.



CONTENTS

Acronyms	xi
Summary	xiii
Samenvatting	xv
1 Introduction	1
1.1 Background	3
1.1.1 The user's role in IoT security and privacy	4
1.1.2 IoT device manufacturers' perspectives and practices	6
1.1.3 Regulatory developments in IoT security.	7
1.2 Research gaps.	10
1.3 Research Objective and Question	11
2 User Expectations	15
2.1 Introduction	16
2.2 Background and Related Work	17
2.2.1 Expectations of IoT security and privacy.	17
2.2.2 Reasonable expectations in law	18
2.2.3 The role of other stakeholders	18
2.3 Method	19
2.3.1 Measuring consumer expectations: the ideal and the learned	20
2.3.2 Survey procedure	20
2.3.3 Vignette design	21
2.3.4 Choice of vignette factors	22
2.3.5 Ethics	24
2.3.6 Participants	24
2.3.7 Vignette and response quality	24
2.3.8 Data analysis.	25
2.4 Results	26
2.4.1 Expectations of manufacturers.	26
2.4.2 The user's role	29
2.4.3 Personal experiences.	31
2.5 Discussion	32
2.5.1 Recommendations.	35
2.5.2 Limitations.	35
2.6 Conclusion	36

3	CRA Expectations	37
3.1	Introduction	38
3.2	Background and Related Work	39
3.2.1	EU support duration regulations.	40
3.2.2	Consumer expectations and perspectives of IoT security and product lifespans	41
3.3	Method	42
3.3.1	Survey Design	42
3.3.2	Recruitment and participants	44
3.3.3	Survey procedure	46
3.3.4	Data analysis.	47
3.4	Results	48
3.4.1	Individual use of smart devices	48
3.4.2	Lifetime expectations of smart devices.	51
3.4.3	Perceptions of software updates and security over devices' lifetime	53
3.5	Discussion	58
3.5.1	Device lifetimes	58
3.5.2	Software Updates and Security over devices' lifespan	59
3.5.3	Recommendations.	60
3.5.4	Limitations.	61
3.6	Conclusion	62
4	Patching Medical IoT Devices	63
4.1	Introduction	64
4.2	Background and Related Work	65
4.2.1	IT Security in the medical domain	65
4.2.2	Patching practices	66
4.2.3	Regulatory landscape of medical IoT devices	66
4.3	Methodology	67
4.3.1	Recruitment and Participants	68
4.3.2	Study Design.	68
4.3.3	Interview Procedure	70
4.3.4	Data Analysis	71
4.3.5	Ethics and Data Protection.	71
4.4	Results	71
4.4.1	Infrastructure	72
4.4.2	Patching pathways.	74
4.4.3	Patching Connected Medical Devices	76
4.4.4	Manufacturer perspective	80
4.5	Discussion	82
4.5.1	Update practices for medical devices	82
4.5.2	Challenges	83
4.5.3	Limitations.	85
4.5.4	Recommendations.	85
4.6	Conclusion	86

5 Conclusion	87
5.1 Empirical Findings	87
5.2 Discussion	89
5.3 Governance Implications	92
5.4 Future work.	98
Bibliography	101
A Appendix for Chapter 2	127
A.1 Survey Instrument	127
A.2 Regression Table	129
B Appendix for Chapter 3	131
B.1 Survey Instrument	132
B.2 Regression Table	137
B.3 Open text response codes	138
B.4 Codebook.	140
C Appendix for Chapter 4	143
C.1 Interview protocol – HDOs	144
C.2 Interview protocol – Manufacturers.	146
C.3 Device Patching Instances	148
C.4 Codebooks	151
Acknowledgements	155
Authorship Contributions	157
List of Publications	159
Datasets	161
About the Author	163



ACRONYMS

ADCO	Administrative Cooperation Group
CERT	Computer Emergency Response Team
CCPA	California Consumer Privacy Act
CRA	Cyber Resilience Act
DDoS	Distributed Denial of Service Attack
EMR	Electronic Medical Records
EAM	Enterprise Asset Management
ENISA	European Union Agency for Cybersecurity
FDA	Food and Drug Administration
GDPR	General Data Protection Regulation
GPS	Global Positioning System
HDO	Healthcare Delivery Organization
ICT	Information and Communication Technology
IT	Information Technology
IVDR	In Vitro Diagnostics Regulation
M-ICT	Medical Information and Communication Technology
MDR	Medical Device Regulation
MAC	Media-Access-Control
MHRA	Medicines and Healthcare Products Regulatory Agency
MR	Magnetic Resonance Imaging
NHS	National Health Service
NIST	National Institute of Standards and Technology
OS	Operating System
PLD	Product Liability Directive
PSTI	Product Security and Telecommunications Infrastructure
RED	Radio Equipment Directive
SaaSMD	Software-as-a-Medical-Device
SBOM	Software-Bill-of-Materials

SG Sale of Goods Directive

UI User Interface

UX User Experience

SUMMARY

The rapid expansion of the Internet of Things (IoT) has increasingly computerized physical objects and tools. Products such as household appliances, cars, industrial machinery, and medical devices are now equipped with sensors, software, and network connections that enable new forms of automation and data exchange. While these developments promise convenience and efficiency, they also introduce new risks. Weak security controls, insufficient privacy safeguards, and inadequate post-market support have repeatedly led to breaches, surveillance incidents, and large-scale cyberattacks. As IoT systems proliferate, the consequences of these vulnerabilities extend beyond individual users or companies to critical infrastructures and society as a whole.

Growing security challenges of the IoT reflect a structural imbalance in the market. Consumers typically lack the information or technical capacity to evaluate or influence a product's security, while manufacturers face little incentive to prioritize it over features or cost efficiency. To address this, governments, particularly within the European Union, are increasingly introducing legislation that codifies how security should be built, maintained, and enforced across the IoT ecosystem. Key among these initiatives are the Cyber Resilience Act (CRA) and the revised Product Liability Directive (PLD), which place explicit emphasis on the expectations of users as a benchmark for determining compliance and responsibility.

The concept of reasonable user expectations has therefore become central to the regulation of IoT products. It provides a flexible legal standard to assess what users can justifiably anticipate regarding the safety and security of their devices. However, despite its prominence in emerging laws, there is no agreed-upon method for determining what these expectations actually are. Courts may consider factors such as prevailing industry practices or product marketing, but empirical evidence of what users themselves expect in practice has been scarce. This creates uncertainty for regulators and manufacturers alike, who must interpret and act on these expectations long before any case law emerges. Against this backdrop, the overarching research question guiding the work is: *What are users' expectations regarding preventive and reactive security measures of IoT devices?*

To answer this research question, this dissertation investigates user expectations at different stages of the IoT device lifecycle: when security or privacy incidents occur, how they are prevented over the device's lifespan, and when devices are used in organizational environments. The studies link these expectations to the broader regulatory concepts of product liability and product conformity, providing evidence that can inform both policy and industry practice.

Chapter 2 focuses on expectations in the context of reactive security, specifically what users expect to happen when an incident occurs with a consumer IoT device. Using a large-scale survey with systematically varied scenarios, it examines how expectations differ across device types, incident kinds, and manufacturer and user responses. The

findings reveal that consumers generally view manufacturers as responsible for resolving security incidents in some capacity, but also display uncertainty if manufacturers would actually respond to privacy-related incidents, and what the user should ideally do. Importantly, the study distinguishes between normative expectations (what users believe should happen) and reasonable expectations (what they think is realistic), providing empirical grounding for the legal interpretation of reasonableness in product liability cases.

Chapter 3 shifts attention to preventive security by exploring users' expectations for how long IoT devices are used and will receive security updates. These insights are directly relevant to the CRA, which ties security support obligations to the expected product lifetime. Drawing on survey data from multiple European countries, it analyzes consumers' expectations of IoT device lifetimes and actual use times, and how these expectations are influenced by the device's nature and individual differences. The results show that users typically expect longer lifespans and support durations than what manufacturers currently provide, revealing a misalignment between market practices and societal expectations.

Chapter 4 extends the analysis to the organizational domain by examining connected medical devices, a critical and highly regulated subset of the IoT ecosystem. Through semi-structured interviews with medical device technicians, medical physicists, and device manufacturers, the study explores how these devices are maintained and patched in practice and what stakeholders' expectations are regarding securing these devices. It highlights a complex coordination between stakeholders, the technical and regulatory barriers to timely updates, and the costliness of ongoing security maintenance. Even in this regulated environment, the study finds persistent challenges in aligning stakeholder expectations and responsibilities, illustrating how systemic and organizational factors shape the feasibility of maintaining security over time.

Taken together, the dissertation studies users' expectations regarding reactive and preventative IoT security, which is based in a relatively young and quickly evolving market. As such, user expectations in this domain, which have legal relevance in upcoming EU legislation, remain empirically unclear. By studying consumers, hospitals deploying medical IoT devices, and medical device manufacturers, we find that expectations of IoT security are varied, context-dependent, and frequently exceed the measures currently in place. As securing the diverse IoT ecosystem is a governance feat for a range of different actors, such as regulators, IoT manufacturers, organizations deploying them, and users, we derive implications for them to support ongoing and future governance efforts. As expectations are heterogeneous but still crucial as a general benchmark for legislation, the findings underscore the importance of grounding legal concepts such as "reasonable expectations" in empirical evidence, providing a bridge between user perspectives, manufacturer practices, and regulatory aims. By clarifying how security is understood and expected in real-world contexts, this work contributes to more effective governance and a more trustworthy foundation for the connected technologies that increasingly shape modern life.

SAMENVATTING

Door de snelle groei van het internet der dingen (IoT) worden fysieke objecten en tools steeds meer gecomputeerd. Producten zoals huishoudelijke apparaten, auto's, industriële machines en medische apparatuur zijn nu uitgerust met sensoren, software en netwerkverbindingen die nieuwe vormen van automatisering en data-uitwisseling mogelijk maken. Hoewel deze ontwikkelingen gemak en efficiëntie beloven, brengen ze ook nieuwe risico's met zich mee. Zwakke beveiligingscontroles, onvoldoende privacywaarborgen en ontoereikende ondersteuning na het in de handel brengen hebben herhaaldelijk geleid tot inbreuken, surveillance-incidenten en grootschalige cyberaanvallen. Omdat IoT-systemen zich steeds verder verspreiden, reiken de gevolgen van deze kwetsbaarheden verder dan individuele gebruikers of bedrijven en hebben ze ook gevolgen voor kritieke infrastructuren en de samenleving als geheel.

De toenemende beveiligingsuitdagingen van het IoT weerspiegelen een structurele onbalans in de markt. Consumenten beschikken doorgaans niet over de informatie of technische capaciteiten om de beveiliging van een product te beoordelen of te beïnvloeden, terwijl fabrikanten weinig stimulans hebben om hieraan voorrang te geven boven functies of kostenefficiëntie. Om dit aan te pakken, voeren overheden, met name binnen de Europese Unie, steeds vaker wetgeving in die voorschrijft hoe de beveiliging in het IoT-ecosysteem moet worden opgebouwd, onderhouden en gehandhaafd. Belangrijke initiatieven in dit verband zijn de Cyber Resilience Act (CRA) en de gewijzigde Product Liability Directive (PLD), waarin expliciet de nadruk wordt gelegd op de verwachtingen van gebruikers als maatstaf voor het bepalen van naleving en verantwoordelijkheid.

Het concept van redelijke verwachtingen van gebruikers is daarom centraal komen te staan in de regelgeving voor IoT-producten. Het biedt een flexibele wettelijke norm om te beoordelen wat gebruikers terecht kunnen verwachten met betrekking tot de veiligheid en beveiliging van hun apparaten. Ondanks de prominente plaats in nieuwe wetgeving bestaat er echter geen overeengekomen methode om te bepalen wat deze verwachtingen precies zijn. Rechtbanken kunnen rekening houden met factoren zoals gangbare praktijken in de sector of productmarketing, maar empirisch bewijs van wat gebruikers zelf in de praktijk verwachten, is schaars. Dit zorgt voor onzekerheid bij zowel regelgevers als fabrikanten, die deze verwachtingen moeten interpreteren en ernaar moeten handelen lang voordat er jurisprudentie ontstaat. Tegen deze achtergrond is de overkoepelende onderzoeksvraag die als leidraad voor het werk dient: *Wat zijn de verwachtingen van gebruikers met betrekking tot preventieve en reactieve beveiligingsmaatregelen voor IoT-apparaten?*

Om deze onderzoeksvraag te beantwoorden, onderzoekt dit proefschrift de verwachtingen van gebruikers in verschillende stadia van de levenscyclus van IoT-apparaten: wanneer er beveiligings- of privacyincidenten plaatsvinden, hoe deze gedurende de levensduur van het apparaat worden voorkomen en wanneer apparaten in organisatorische omgevingen worden gebruikt. De studies koppelen deze verwachtingen aan de bredere

regelgevingsconcepten van productaansprakelijkheid en productconformiteit, en leveren bewijs dat zowel het beleid als de praktijk in de sector kan informeren.

Hoofdstuk 2 richt zich op verwachtingen in de context van reactieve beveiliging, met name wat gebruikers verwachten dat er gebeurt wanneer zich een incident voordoet met een consumenten-IoT-apparaat. Aan de hand van een grootschalige enquête met systematisch gevarieerde scenario's wordt onderzocht hoe de verwachtingen verschillen naargelang het type apparaat, het soort incident en de reacties van fabrikanten en gebruikers. Uit de resultaten blijkt dat consumenten over het algemeen vinden dat fabrikanten in zekere mate verantwoordelijk zijn voor het oplossen van beveiligingsincidenten, maar ook onzeker zijn of fabrikanten daadwerkelijk zouden reageren op privacygerelateerde incidenten en wat de gebruiker idealiter zou moeten doen. Belangrijk is dat de studie onderscheid maakt tussen normatieve verwachtingen (wat gebruikers denken dat er zou moeten gebeuren) en redelijke verwachtingen (wat zij realistisch vinden), waardoor een empirische basis wordt geboden voor de juridische interpretatie van redelijkheid in productaansprakelijkheidszaken.

Hoofdstuk 3 verschuift de aandacht naar preventieve beveiliging door te onderzoeken wat de verwachtingen van gebruikers zijn met betrekking tot de gebruiksduur van IoT-apparaten en de duur van beveiligingsupdates. Deze inzichten zijn direct relevant voor de CRA, die beveiligingsondersteuningsverplichtingen koppelt aan de verwachte levensduur van het product. Op basis van enquêtegegevens uit meerdere Europese landen analyseert het de verwachtingen van consumenten ten aanzien van de levensduur van IoT-apparaten en de daadwerkelijke gebruiksduur, en hoe deze verwachtingen worden beïnvloed door de aard van het apparaat en individuele verschillen. Uit de resultaten blijkt dat gebruikers doorgaans een langere levensduur en ondersteuningsduur verwachten dan wat fabrikanten momenteel bieden, wat wijst op een discrepantie tussen marktpraktijken en maatschappelijke verwachtingen.

Hoofdstuk 4 breidt de analyse uit naar het organisatorische domein door gekoppelde medische apparaten te onderzoeken, een kritische en sterk gereguleerde subgroep van het IoT-ecosysteem. Aan de hand van semi-gestructureerde interviews met technici van medische apparatuur, medisch fysici en fabrikanten van apparatuur onderzoekt de studie hoe deze apparaten in de praktijk worden onderhouden en gepatcht en wat de verwachtingen van belanghebbenden zijn met betrekking tot de beveiliging van deze apparaten. Het benadrukt de complexe coördinatie tussen belanghebbenden, de technische en regelgevende belemmeringen voor tijdige updates en de hoge kosten van voortdurend onderhoud van de beveiliging. Zelfs in deze gereguleerde omgeving constateert het onderzoek aanhoudende uitdagingen bij het afstemmen van de verwachtingen en verantwoordelijkheden van belanghebbenden, wat illustreert hoe systemische en organisatorische factoren de haalbaarheid van het handhaven van beveiliging in de loop van de tijd bepalen.

Al met al onderzoekt het proefschrift de verwachtingen van gebruikers met betrekking tot reactieve en preventieve IoT-beveiliging, die is gebaseerd op een relatief jonge en snel veranderende markt. Als zodanig blijven de verwachtingen van gebruikers op dit gebied, die juridisch relevant zijn in de komende EU-wetgeving, empirisch onduidelijk. Door onderzoek te doen naar consumenten, ziekenhuizen die medische IoT-apparaten gebruiken en fabrikanten van medische apparatuur, hebben we geconstateerd dat de

verwachtingen ten aanzien van IoT-beveiliging uiteenlopen, afhankelijk zijn van de context en vaak verder gaan dan de maatregelen die momenteel worden genomen. Aangezien het beveiligen van het diverse IoT-ecosysteem een governance-taak is voor een reeks verschillende actoren, zoals regelgevers, IoT-fabrikanten, organisaties die deze apparaten gebruiken en gebruikers, leiden we hieruit implicaties af voor hen om lopende en toekomstige governance-inspanningen te ondersteunen. Aangezien de verwachtingen heterogeen zijn, maar toch cruciaal als algemene benchmark voor wetgeving, onderstrepen de bevindingen het belang van het baseren van juridische concepten zoals “redelijke verwachtingen” op empirisch bewijs, waardoor een brug wordt geslagen tussen het perspectief van de gebruiker, de praktijken van de fabrikant en de doelstellingen van de regelgever. Door te verduidelijken hoe beveiliging in de praktijk wordt begrepen en verwacht, draagt dit werk bij aan effectiever bestuur en een betrouwbaarder fundament voor de verbonden technologieën die het moderne leven in toenemende mate vormgeven.



1

INTRODUCTION

A growing number of physical objects around us are becoming increasingly network-connected. With smaller and more advanced hardware components such as Microcontroller Units, sensors, and actuators, products like home appliances, medical devices, cars, or industrial machinery can be connected to other devices and networks, such as the Internet. This enables functionalities not possible before, such as remote monitoring, predictive maintenance, real-time data analytics, and automated control. Together, these interconnected systems form the Internet of Things (IoT), “*the network of devices that contain the hardware, software, firmware, and actuators which allow the devices to connect, interact, and freely exchange data and information*” [171]. The IoT ecosystem is proliferating rapidly, and it is expected to follow this trend in the upcoming years: it has been estimated that by 2030, 40 billion IoT devices will be online [214], with an estimated annual growth in market volume of 10% [216].

This increasing connectivity of *things* also introduces novel risks for users, organizations, and society at large. Security and privacy threats have been an inherent part of the realm of computers, but now also apply to a larger number of products used ubiquitously and running critical infrastructure, as they increasingly become computers. However, security measures for IoT devices have been repeatedly reported to be lacking, such as insufficient access control [131], lacking encryption [175], or vulnerable companion applications on IoT users’ smartphones [163]. As the IoT market is driven by fast innovation cycles and short development times, security and privacy have frequently been reported to be an afterthought during development, as features and usability are prioritized [47, 126]. Crucially, to keep IoT devices secure, risks need to be managed and mitigated where needed well after the product has been put onto the market, as once devices are not supported in this way and are considered “end of life”, they can quickly become vulnerable [253].

Insufficient security and data protection measures in IoT devices led to real harm, such as users being spied upon [55, 201, 252], abuse and harassment in romantic relationships [28, 222], financial crimes [215], or IoT devices being used as entry-points into organizational networks[41, 118, 254]. A further malicious application of IoT devices

commonly observed in the wild is the formation of botnets consisting of thousands of infected IoT devices to run distributed denial of service attacks (DDoS) [42, 126, 207]. Due to the large volume of devices and their simultaneous requests to the target system, botnets can overwhelm network infrastructure, such as websites of public services, thus rendering them inaccessible for legitimate visitors. Recent reports highlight the significant role of compromised consumer IoT devices in DDoS attacks with record-breaking volume [129, 261]. Thus, the sheer number of actively connected IoT devices that can become vulnerable to such malware over their lifetime poses a significant risk to the network infrastructure that our society relies on.

It has been argued that lacking IoT security signals a market failure [31, 204, 250], where users are left in a poor position to manage security for themselves, and manufacturers lack incentives to develop more secure products. Users usually face information asymmetry, where they cannot know what security or data collection properties a given IoT device has, a lack of usable device settings to be able to opt for secure or privacy-preserving options, or no way of diagnosing if a device has been attacked or how to fix it [27]. Manufacturers, on the other hand, have been lacking incentives to incorporate stronger security into their products, as well as guidelines and regulations that describe how to reach a certain level of security for IoT devices. As a result, security has often been termed an afterthought during product development [47, 205], and the assumption of informed consumer choice during purchase cannot hold if they do not have access to the relevant information on the product's security properties. Thus, externalities in the form of societal security risks are likely to persist in this market dynamic, as manufacturers lack incentives to design more secure products and users lack the information and tools to manage security independently, if needed.

In cases of market failures, government interventions are commonly called for [200, 204], usually involving the design and enforcement of legislation. However, legal instruments to define and enforce requirements for IoT device security have generally been lacking so far. In the related domain of product safety, two legal concepts are used to prevent or navigate products failures and define legal responsibilities: product *conformity*, which refers to the product meeting a set of minimum legal or technical requirements to prevent harm; and product *liability*, which defines the legal responsibility of the involved parties (usually the manufacturer, seller, and/or user) for harm caused by a defective product after it has been placed on the market. Applied to IoT security, this would mean that an IoT device needs to comply with a certain set of security requirements to *conform* to the law, and that there is a legal process in place to assess and enforce who is *liable* for any damage caused by poor security. However, traditional product conformity and liability regimes do not translate well to security, as they have been built around the idea of a physical good whose safety characteristics can be assessed once, at the point of manufacture, against a relatively stable set of hazards. In security, however, vulnerabilities, threats, and the technical "state of the art" are ever-changing and require ongoing post-market monitoring. Furthermore, harm in this context is often more intangible and difficult to measure than physical harm resulting from safety defects. Consequently, neither regulators nor users, be it consumers or organizations, have robust assurances or tangible legal or technical levers to demand or enforce meaningful security in IoT devices.

To address this state of the market, policy initiatives aiming at regulating IoT device security have accelerated rapidly in recent years. One of the strongest regulatory pushes comes from the European Union, which is implementing legislation specifically targeting product conformity [230, 243] and liability [242] for security. Due to the complex and dynamic nature of this domain, legislation needs to be flexible enough to account for the contextual intricacies of a given case and avoid requirements that might be too rigid for IoT manufacturers to implement effectively. This legal degree of freedom often manifests in legislative texts as a reference to users' or consumers' "reasonable" expectations [61, 91, 195]. In these cases, the expectations of the public at large, such as those of users, serve as a baseline for legal interpretation in the courts to determine what is "reasonable" in the given context. There is no standardized methodology provided for determining this baseline, and it is defined by the courts for the specific case at hand.

User expectations are also a crucial component of emerging product security legislation. The Cyber Resilience Act refers to the expectations of users as a benchmark for what is a reasonable post-market support duration to uphold products' conformity to the law's requirements [243], and the Product Liability Directive takes expectations as a benchmark for what security a user can be entitled to expect in determining liability [242]. However, as the reasonableness of an expectation is defined by the courts on a case-by-case basis, this implies that cases actually reach a court for a judge to decide on what constitutes a reasonable or unreasonable expectation. This leads to uncertainty for actors who must make decisions based on this abstract legal notion, as relevant precedents from case law may be lacking, and going to court may constitute a legal and financial risk. For instance, manufacturers of IoT devices and market surveillance authorities enforcing the recent and upcoming regulations require a tangible and joint understanding of users' expectations regarding IoT security, so they can consider this in their decision-making during product design, compliance activities, and, in the case of authorities, enforcement of the law. It further leads to persistent uncertainty for users of IoT devices regarding their rights and obligations, as their "expectations" could mean many things.

In sum, the market for IoT devices is currently undergoing a transitional phase towards stricter security regulations. However, the key concept of user expectations regarding security remains uncertain in this complex market condition.

1.1. BACKGROUND

This section presents relevant academic work on the past and current state of security in the IoT market, with a focus on users and manufacturers. It then delves into interventions and initiatives from governance to tackle market failures in this domain. Section 1.1.1 summarizes empirical research about the position users have been in regarding IoT security. It examines their perspectives, lived realities, and expectations to better understand the user's assumed and de facto role in securing IoT devices. Section 1.1.2 then examines research on manufacturer practices and perspectives regarding IoT product security to highlight internal processes, incentive structures, and market dynamics under which they are operating. Finally, Section 1.1.3 presents ongoing legislative developments related to IoT security and privacy, and how the concept of consumer expectations is woven into it.

1.1.1. THE USER'S ROLE IN IOT SECURITY AND PRIVACY

Users of IoT devices commonly face a multitude of challenges when purchasing, configuring, using, and disposing of these products, as well as navigating potential security and privacy risks and incidents. This contributes to a persistent state of IoT vulnerability, as secure usage and configuration, like regularly installing updates or changing default passwords, can have a direct effect on devices' level of security. Indeed, a large portion of the responsibility for ensuring security throughout the devices' lifecycle has been predominantly with the user so far, as they have been assumed to actively engage in such behaviors, often without being provided with usable tools and information on how to do so.

During purchase, users usually face inconsistent or incomplete information about security and privacy-related product characteristics, such as how and for how long it will receive security updates from the vendor, which security measures are implemented in the device, or how device data flows and is used [25, 81, 104, 196, 198]. Thus, the consumer IoT market was coined a "market for lemons", as product complexity and incomplete information lead to an information asymmetry between buyer and seller [90]. While disclosure of IoT product security is thus often piecemeal, prior research suggests that consumers exhibit a preference for stronger security and privacy measures in IoT devices [73, 160, 164], and are willing to pay a premium for this [24, 72, 95, 250]. An approach commonly proposed to combat information asymmetry during purchase is IoT product security labels [70, 73, 124, 160]. Labels can provide potential buyers with a summary of key privacy and security attributes, either in binary (e.g., seals of approval), graded (e.g., ordinal scales reflecting the degree of security), or descriptive (listing and describing security properties) formats. IoT security and privacy labels are actively deployed or planned in several countries, such as Singapore [180], the UK [109], Germany [83], and the USA [80]. While consumers seem willing to pay more for devices with better security, thus being an incentive for manufacturers to signal this with a label, several challenges in rolling out product labels have been reported. For instance, doubts if a label can accurately reflect the security characteristics of a device [89] and difficulties for end users in interpreting the varying technical label attributes, like misinterpreting the end of update support given on the label as an expiry date of the product [21, 70, 109]. Furthermore, labeling schemes are largely voluntary for manufacturers to add to their products (as in, e.g., the USA or Germany), and their effectiveness in swaying consumers towards more secure options remains empirically unclear.

Also during operation, that is, after IoT devices have been purchased, users commonly face limited insight into and control over their devices' security and privacy properties. This is despite repeated empirical findings that their concerns regarding security risks are high [44, 105, 262]. A lack of control is often caused by missing or limited user interfaces and missing or difficult to find security and privacy settings [44], the "notice and choice" paradigm of privacy policies [147], and inconsistent or misleading advice on how to manage security and privacy of IoT devices [16, 25, 244, 245]. One line of research proposes to improve insights and control for home users with dashboards [4, 39, 137], where one user interface summarizes connected consumer IoT devices and visualizes data flows. However, this approach assumes users' willingness and ability to actively monitor and manage all their IoT devices, rendering them their own system administra-

tor, which arguably moves the burden of ongoing risk monitoring to them. As prior work also reports on gaps in users' mental models about their devices' security [5, 262], this approach seems unfeasible, especially if the number of connected devices in a household rises.

Installing the latest security updates is generally considered best practice for users to mitigate risks, as also widely advised by governments to consumers [244, 245]. Indeed, installing patches often constitutes one of the only available controls for users to actively approach security [104]. However, empirical work suggests that a plethora of challenges and uncertainties makes this difficult for users, such as lacking clarity regarding updates' purpose and methods of delivery [104], common bundling of security updates with feature updates, which introduces the tradeoff for users between having to accept undesired functional changes to the device in order to receive the latest security update, or users' uncertainty regarding the implications of the end of update support and if this would lead to increased security risks [70, 102, 103]. Thus, significant challenges also exist for this key preventive user security behavior.

Beyond prevention, options for remediating from security and privacy incidents are limited and difficult to predict for users. When faced with potential security and privacy incidents, there is no clear path to resolution. Detecting that a device is compromised or data is shared illicitly is difficult, if not impossible, for users, as there are usually no direct feedback or diagnostic tools available [201, 224]. Even when directly notified, prior work highlights that it is difficult for users to implement the necessary and sometimes technically complex steps to remediate [27]. Thus, in the face of potential security risks or incidents, users often seek informal advice from others [10, 111, 173], attempt ad-hoc technical solutions themselves [87, 105], stop using devices altogether [27, 201, 224], or simply continue usage due to a perceived low risk [123, 262]. Partly due to this lack of options and dependency on device manufacturers, prior research has also documented a general resignation among users when it comes to security and privacy of their devices [101, 210]. When reaching the end of the device use time, users are furthermore prompted to remove their usage data from the device to avoid leakage of potentially sensitive information, such as credentials like passwords. However, prior research found that IoT devices generally hold more data than users expect, with only half of users reporting removing data before disposal. Moreover, the majority of data on devices studied was saved in plain text, and disposal methods like overwriting and device resetting often did not ensure effective data removal [140].

In sum, the current status quo largely requires users of IoT devices to manage their devices' security and personal information to an unrealistic extent. Actual tools and options provided to them to do so are limited across the entire device lifespan: from an informed purchase decision, to secure usage, and up to the device's end of life and data disposal. However, when being queried directly, users see the responsibility for securing smart devices as being shared across themselves, device manufacturers, and the government [101]. Thus, a gap appears to exist between users' expectations and the current state of IoT device security, as well as a lack of clear distribution of legal responsibilities among stakeholders to ensure effective risk mitigation.

1.1.2. IOT DEVICE MANUFACTURERS' PERSPECTIVES AND PRACTICES

To understand the current state of security in the IoT market in the context of looming regulations, an inside perspective on IoT manufacturers is required to understand processes, priorities, and challenges when designing, producing, selling, and maintaining their products. Previous research examines how manufacturers integrate security and privacy into their products during development, revealing a complex, multi-layered process. Previous work has studied "security by design": e.g., by supporting developers with more usable tools and feedback mechanisms to implement security solutions into products, as developer tools have been frequently reported to have shortcomings and implementation challenges [9, 221, 256]. While easing the way for developers to follow more secure practices is an important step, other challenges for manufacturers have been reported that seem more systemic in nature and thus require different approaches than improving tooling.

For instance, previous empirical work with IoT manufacturers noted that implementing security incurs considerable costs for the organization. This included adopting more secure hardware components [227], maintaining security of devices over time (e.g., with security patches) [45], or certifying products according to security standards [145]. Thus, security is often subject to internal budget negotiations across different departments, such as security teams, product teams, or management. As successful security outcomes are difficult to quantify, such budget discussions were reported to be non-trivial during development for security teams [227], thus requiring active backing from management and a security-focused organizational culture [169]. Prior work also found that teams that could collaborate and contribute to more usable security for IoT users, that is, User Experience (UX) Design and product security teams, did not usually do so, as UX and security were considered entirely unrelated and did not share compatible processes [43]. Further challenges in incorporating security more thoroughly in development was reported a lack of expert personnel with the relevant skills [145, 227] and the process of integrating hard- and software components from downstream suppliers, as IoT manufacturers would have to rely on these components' respective interface's compatibility and usability to integrate into the product together with components from other vendors [227] as well as their timely fixing of vulnerabilities by the original supplier [45].

A further major theme from IoT manufacturer-focused research is the role of standardization and regulation. Standards and regulatory compliance were one of the dominant drivers for manufacturers to implement security into products in sectors and situations where legislation was already in place. For connected medical devices, compliance was one of the main drivers for developers to follow security practices [255], as was the case in the automotive and financial sectors, where standards and legislation such as UNECE Regulation No. 155 [82], and GDPR [231] defined security and privacy processes during development and maintenance [227]. Furthermore, Chalhoub et al. [43] found that for smart security camera manufacturers, a sector with substantially less regulatory pressure than medical devices or cars, User Experience designers worked together with legal departments to provide users with better interfaces for data sharing options in the context of GDPR. As seen earlier, security and UX teams did not collaborate to deliver intuitive interfaces for security features, thereby highlighting the potential impact of legislation (in this case, GDPR) on the design of user-facing technology.

While regulations with mandatory requirements are a strong driver to force manufacturers to incorporate security and privacy more thoroughly into their products, a lack of policy, specifically in the form of inconsistent or vague standardization, is commonly reported to be a major challenge for the industry. Practitioners highlighted a lack of consistent and actionable standards for securing IoT devices, as existing standards were often vague, not providing an actionable process for doing this, or even lacked a clear conceptualization and definition of IoT security, thus contributing to a "wild west" scenario, where practices were ad-hoc and highly heterogeneous [169, 227]. Furthermore, this makes it difficult for manufacturers to find a shared language with their customers, other businesses, or consumers, to convince them that their security is "good enough". However, Mandal et al. [145] report that while standards might be fragmented and vague, their absence actually led to a neglect of security during development, thus highlighting how the mere existence of security standards can improve it due to developers considering it earlier in development. In the same study, practitioners also highlighted how the lack of any mandatory regulation behind a standard disincentivizes developers to use them. Thus, industry professionals [45, 145, 227], authors of the previously presented work [45], and policymakers [245] have repeatedly called for more comprehensive, concrete, and legally binding regulations and guidelines to provide clear requirements for IoT products and predictable and actionable processes for manufacturers to implement security.

1.1.3. REGULATORY DEVELOPMENTS IN IoT SECURITY

As sections 1.1.1 and 1.1.2 have outlined, the market for IoT devices has been plagued by a plethora of security challenges, including transparency issues for users and a lack of incentives and workable guidelines for manufacturers. Further complicating matters, existing laws governing the security of IoT products have been lacking. From a preventative (or ex-ante) perspective, *product conformity* legislation was missing to bring IoT security to a certain baseline level by defining legally binding requirements that these systems need to comply with in order to be sold, so that the likelihood of adverse outcomes is reduced. From a reactive (ex-post) perspective, *product liability* laws and schemes were missing to define legal clarity for responsibilities in case (vulnerable) IoT devices were involved in causing harm. Thus, if IoT devices were attacked or personal user data compromised, there was often no legal clarity on which actors were responsible for prevention and remediation, and to what extent.

However, policy responses to regulate the IoT device market more thoroughly have been developing rapidly to catch up. The last years mark a substantial acceleration of cybersecurity governance in the IoT space, with dozens of new regulations and standards being written, implemented, and enforced. This growing response from policymakers translates into legislative initiatives that define legally binding obligations, going beyond previous voluntary approaches, such as product labels, optional certification, or industry guidelines. While legislation in this space develops dynamically in many regions, like Asia [98, 158, 180] and North America (e.g., the U.S. Cyber Trust Mark [80] or the IoT Cybersecurity Improvement Act [1]), this section will focus predominantly on the European region (including the UK), as one of the most active and influential regions globally in terms of regulation [29].

For IoT devices in general, the EU Radio Equipment Directive (RED), a directive aimed at all products with communication capabilities via radio (e.g., Wi-Fi, Bluetooth, cellular) was extended via a Delegated Act in 2022 to add essential cybersecurity requirements for IoT devices, such that they must not harm communication networks or misuse network resources, include safeguards for personal data and user privacy, and resist fraud [230]. Manufacturers will need to implement measures to ensure these requirements for products placed on the European market by August 2025 to receive the RED CE mark and be able to sell their devices in the European Union. Similarly, the UK has passed the Product Security and Telecommunications Infrastructure (PSTI) Act in 2022 [229]. It has been in force since 2024 and requires manufacturers, importers, and/or distributors of "consumer connectable products" to meet minimum security requirements, such as not integrating default or easy-to-guess passwords or disclosing information on security update support duration. Manufacturers must self-attest that their products comply with these requirements and provide the required documentation, which must be verified by sellers and distributors in the UK. Furthermore, in both the EU and the UK, the General Data Protection Regulation (GDPR) applies to any IoT device or service handling personal data [231]¹. It imposes rules on data processing, consent, and the security of personal data, with IoT manufacturers and service providers having to ensure privacy-by-design in their devices, meaning minimization of and informing users on data collection while also keeping them in control of their data.

More recently, the EU Cyber Resilience Act (CRA) was adopted in 2024 and is scheduled to take effect in 2027 [243]. It applies to all "products with digital elements" and is arguably the largest security-related product legislation in the EU to date. It marks a significant shift by establishing mandatory, horizontal cybersecurity requirements for all products with digital elements placed on the EU market, including most IoT devices and components thereof. The CRA's EU-wide horizontal nature contrasts with directives such as the RED, which require transposition into national law. The CRA introduces a range of obligations for manufacturers, who will have to ensure security is incorporated into IoT devices' design, production, operation, and maintenance ("security-by-design"), with potential security risks to be monitored and mitigated (e.g., with security updates). Crucially, vulnerabilities will need to be monitored and mitigated if needed throughout the entire expected lifetime of the product, thus introducing a strong focus on post-market surveillance and on the upkeep of security measures *over time*. When defining the expected lifetime of a product type, manufacturers must consider the expectations of users, as well as the nature of the product and any relevant union laws. The CRA thus raises accountability and responsibility for the manufacturers of IoT devices and all products with digital elements, by directly connecting the lifetime dimension with society's (i.e., users') expectations. To ensure practicable and streamlined processes for manufacturers, the CRA will be supplemented by a range of standards [49], as well as the Cybersecurity Act [234], which introduces an EU-wide horizontal certification scheme for all products within the EU [50].

While the CRA, RED, and PSTI primarily pertain to product conformity, i.e., the minimum requirements a product must fulfill, liability resulting from security or privacy inci-

¹With the UK currently using the UK-GDPR as well as the Data Protection Act 2018 <https://www.gov.uk/data-protection>

dents is now also increasingly targeted by legislation. The EU Product Liability Directive (PLD)[242], originating in 1985, has been updated and subsequently adopted in 2024 to account for liability for modern technologies. Software, digital services, and AI-enabled products are now also included in the definition of a "product". The Directive's definition of *damage* now also includes the destruction or corruption of data and damage to psychological health. As successful attacks on vulnerable IoT devices, such as cameras or baby monitors, can have a psychological impact [201], this could now be considered damage under the PLD. Products can furthermore be considered defective where they do not "provide the safety that a person is entitled to expect or that is required under Union or national law", where all circumstances of the product will be considered when determining if it was defective. Crucially, this list of circumstances includes cybersecurity requirements from the regulations previously mentioned, namely the CRA, RED, or GDPR. Manufacturers can also be held liable if the product's defectiveness stems from new updates (i.e., introducing the defectiveness) or vulnerabilities that the manufacturer has failed to patch with "software updates or upgrades necessary to maintain safety". Also, a product can be considered defective by courts if the claimant "faces excessive difficulties, in particular due to technical or scientific complexity, in proving the defectiveness of the product(...)", thus shifting the burden of proof away from, e.g., a user having difficulties to "proof" a software vulnerability in their device led to damages, and making it easier for them to win claims. EU member states must implement the new PLD into their national laws by December 2026; products sold before this date will still fall under the previous liability directive.

Beyond these more horizontal pieces of legislation, sector-specific regulations apply to specific types of IoT devices used in organizational settings, such as industrial machinery, medical devices, or naval vessels. Medical devices are a prominent example of products that are increasingly connected and thus part of the Internet of Things, while being considered critical infrastructure. As such, they are regulated under their own umbrella legislation and are procured and deployed by organizations (usually healthcare delivery organizations) and not by consumers. Medical devices are regulated under the Medical Device Regulation (MDR) or In Vitro Diagnostics Regulation (IVDR), depending on the use case [232, 233]. As summarized in the official EU's medical device guidance on cybersecurity [238], security is considered part of device safety, with connected medical devices required to be designed and manufactured to prevent unauthorized access and manipulation, requiring a continuous risk-management process to monitor and reduce arising risks to patient safety (such as a new software vulnerability). Changes to the system, such as security updates, need to be verified and validated to ensure continuous device performance and safety. Furthermore, hospitals and manufacturers of connected medical devices in the EU are governed under the Network and Information Security Directive (NIS2), which mandates such providers of critical infrastructure to implement security risk management in their infrastructure and to report significant incidents to authorities[240].

Taken together, a large corpus of newly drafted and updated legislation targeting product security highlights the broader shift in EU policy to treat security as a central pillar for the EU's digital future [51].

Crucially, reasonable user or consumer expectations as a concept are used through-

out this legislation. The CRA refers to reasonable user expectations as a benchmark for determining expected use times (and thus, support durations) for product categories [243]. The Product Liability Directive ties the safety (here, security) measures that a product should contain to what "*a person is entitled to expect*". Other examples of product legislation with references to reasonable expectations can be found in the Digital Services and Goods Directive [236] and the General Product Safety Regulation [241], highlighting the widespread use of this concept in EU law to determine which requirements a product should legally fulfill. The concept of reasonable user expectations remains abstract in nature to keep the legislation adaptive and flexible. This has been termed the "open texture" [110] of law to preserve judicial discretion for the courts to handle the complexity of real-world scenarios and avoid being too rigid to account for the sheer heterogeneity of modern digitally-driven products.

1.2. RESEARCH GAPS

Reasonable user expectations are an inherently abstract and context-dependent concept that is defined in the courts on a case-by-case basis. As such, legislation does not provide a clear methodology for assessing it. Geiregat [91] notes that there are many potential sources to define expectations *in court*, ranging from the law itself (e.g., security expectations shaped over time by laws) to what the market offers (i.e., expectations based on similar products or from product marketing). Thus, an empirical lens on user expectations and common security practices can provide insights for implementing security-related regulations in practice, without having to wait for lawsuits and court precedents.

However, empirical work on user expectations in a legal context is scarce. Prior research on user expectations surrounding IoT security has predominantly considered expectations as a normative construct, that is, as the preferences of users [162, 212, 220]. This includes a wide range of different user preferences, such as security [220, 263] and privacy [15, 146] features wanted for IoT devices, security concerns users would want a solution for [44, 105, 262], desired update support durations for IoT devices [160, 164], or which actors should ideally take on responsibility for IoT devices' security and privacy [101]. This line of research typically does not conceptualize users' expectations as a legal construct and therefore does not analyze them within a legal context. While it provides valuable insights into users' wishes and needs for user-centric security solutions and formulates policy recommendations based on empirical results, the studies and measured normative expectations are not explicitly framed as a legal concept and are not directly tied to rights or obligations outlined in legal texts. Thus, previous work studying user expectations largely remains limited in its legal implications, while section 1.1.3 highlights the salient role of user expectations in product conformity and liability legislation for IoT device security. For instance, the studied construct of normative user expectations does not fully capture the legal concept of expectations, as they also have to be "reasonable". A claimant might argue in court that they expected (i.e., wanted) military-level encryption and 30 years of patch support for a low-cost smart light bulb, which is likely a normative expectation, but perhaps not a reasonable one in the eyes of the judge.

Furthermore, the vast amount of prior empirical user research on IoT security has been conducted in the space of consumer products, taking the user or consumer per-

spective (e.g., [72, 103, 144, 162, 202, 215, 220, 263]). However, research that takes the perspective of stakeholders engaged in organizations using IoT products and that explores domains beyond consumer-focused IoT remains sparse. Security-related product legislation also applies to a vast number of IoT products used in organizational settings, such as enterprise IT environments, critical infrastructure, or industrial IoT. Thus, it also relates to stakeholders' expectations within these organizational settings, which operate under different technical, regulatory, and contractual conditions than consumers. For instance, the healthcare sector operates an increasingly larger infrastructure of connected medical devices, and previous literature that positions empirical findings about these devices and their security within the regulatory landscape is lacking, as it mostly focuses on the topic through the lenses of traditional IT (i.e., non-IoT) infrastructure [65–67], security challenges and risks at hospitals in general [7, 52, 69, 100, 120, 151, 247], or legal analysis [141].

In sum, this work addresses the following gaps in the previous literature.

1. We lack empirical insights into the user expectations in case of emerging security incidents for IoT devices and possible responses in the context of product liability.
2. We lack empirical insights into the user expectations of preventative security measures in IoT devices and how long they will be upheld in the context of product conformity.
3. We lack empirical insights into stakeholder expectations and practices regarding IoT device security in organizational settings, such as the healthcare sector.

1.3. RESEARCH OBJECTIVE AND QUESTION

The overarching objective of this dissertation is to empirically assess user expectations of security at different stages of the IoT device lifecycle. With rapidly evolving and young technologies like the Internet of Things and cybersecurity, it remains nebulous what society, whether it be consumers or organizations, expects, let alone what might be "reasonable". The law does not provide a benchmark or methodology to determine this. Thus, this dissertation examines expectations regarding *responding* to emerging security and privacy vulnerabilities and incidents in the context of product liability (*reactive security*, Chapter 2), as well as expectations regarding the ongoing care of IoT devices over their lifespan to *prevent* incidents, that is, in the context of product conformity (*preventative security*, Chapter 3). As "users" not only refers to consumers, but also stakeholders in an organizational capacity who maintain and use them, this work also assesses practices and expectations regarding the reactive and preventive security of connected medical devices over their lifespan in an organizational setting (Chapter 4).

Furthermore, this work aims to empirically examine user behaviors and practices regarding IoT security to relate them to measured expectations, gain insights into common behavior among actors in the current market, and contrast them with legal provisions. Specifically, Chapter 2 also studies users' own experiences with and responses to security incidents, Chapter 3 considers users' IoT device use times and security mitigation actions, and Chapter 4 assesses stakeholders' patching practices at healthcare-delivery

organizations. In this way, the empirical lens sheds light on actual practices and establishes instances of real-world context to understand what might be considered "reasonable" and how current practices relate to the regulatory provisions. In doing so, it aims to inform ongoing discussions and developments in policy and industry regarding IoT security legislation by reducing uncertainty surrounding the concept. Key actors in the market require a workable interpretation and approach today, as waiting for potential precedents from future case law is infeasible. For instance, IoT manufacturers need to plan for compliance with upcoming legislation, such as the CRA, and market surveillance authorities enforcing the law require an actionable understanding of what users expect to use it as a criterion for market oversight.

This dissertation thus pursues the overarching research question:

What are users' expectations regarding preventive and reactive security measures of IoT devices?

This overarching research question is addressed by three different studies presented in the following chapters. Each study focuses on different sub-research questions and addresses them with varying research methods.

STUDY 1: WHEN SECURITY AND PRIVACY FAIL: UNDERSTANDING CONSUMER EXPECTATIONS

Chapter 2 describes a quantitative survey study to understand consumer expectations in the context of potential product liability, that is, in case consumer IoT devices face a security or privacy incident or vulnerability during their usage, such as unauthorized access by malicious attackers (security) or illicit data collection by the manufacturer (privacy), while differentiating between normative and "reasonable" expectations. As liability is context-dependent and determined on a case-by-case basis, the study deploys a survey design that allows for systematically varying "context" by assessing if and how expectations differ for different IoT device categories, security and privacy incidents, and manufacturer and user actions. It thus considers consumer expectations concerning the nature of the compromised device type, incident, and the expected responsibility for *reactive* actions by IoT users and manufacturers. It further reports on experiences with such incidents faced by survey respondents personally.

It is driven by four research questions:

- *RQ1* : What do consumers expect how manufacturers *will* respond to emerging privacy and security risks with IoT devices?;
- *RQ2* : What do consumers expect how manufacturers *should* respond to emerging privacy and security risks with IoT devices?;
- *RQ3* : Do expectations differ across product types and threat events?; and
- *RQ4* : How do participants evaluate the user's responsibility to handle emerging privacy and security risks with IoT devices?

STUDY 2: PREVENTING FAILURES: EXPECTATIONS OF SECURITY SUPPORT OVER DEVICE LIFETIMES

Chapter 3 moves from the *reactive* lens of product liability to *preventative* mitigation in the form of product conformity by studying user expectations regarding the upkeep of

minimum security requirements for consumer IoT devices in the context of the Cyber Resilience Act. As the mandatory duration of security support is correlated with users' expectations about the IoT device's lifespan in the CRA, this study also follows a survey approach in measuring the expected lifespans for a range of different types of consumer IoT devices. It further considers respondents' behaviors with their own devices (like installing updates) and factors underpinning the formation of expectations, like personal experiences or anticipated device usage. As the CRA is horizontal legislation applying to all European Member States, the chapter also reviews survey findings from various EU countries.

The chapter pursues the following Research Questions:

- *RQ1* : How do consumers use their smart devices, and for how long?;
- *RQ2* : How long do consumers expect different smart device categories to last, and which factors influence these expectations?;
- *RQ3* : How do consumers perceive security and software update support over smart devices' lifespans?, and;
- *RQ4* : Are there differences among EU member states regarding consumers' smart device usage, expectations, and security and software support perceptions?

STUDY 3: SECURING CONNECTED MEDICAL DEVICES: PATCHING PRACTICES AND EXPECTATIONS AT ORGANIZATIONS

Chapter 4 then takes an organizational lens on preventative and reactive security in practice by studying how connected medical devices are deployed and secured in their operational environment at Healthcare Delivery Organizations. Connected medical devices are IoT, like the previously studied consumer devices. However, they are operated in a fundamentally different context. They are usually procured in formalized processes by organizations (like hospitals), integrated into the network environment, managed, and maintained by dedicated, professional practitioners (like system administrators and (bio-)medical technicians), used for safety-critical use cases (patient care), and fall under a different regulatory umbrella. Thus, the chapter considers such connected medical devices as a case study for IoT devices as part of an organizational infrastructure, and assesses expectations, practices, and experiences of professional stakeholders involved in securing these devices by conducting semi-structured interviews, specifically with medical technicians, medical physicists, and product security experts at medical device manufacturers.

Medical devices partly fall under other security-related regulations than consumer products². As these medical device regulations have been in effect for a longer period, this chapter also considers connected medical devices as a case study of how regulations impact security practices and how post-market surveillance and risk mitigation, in the form of security updates, work in practice. Finally, the chapter specifically includes the manufacturer's perspective to understand how they navigate regulatory requirements for their products and how security product measures and processes are implemented internally, and what kind of challenges emerge from this or from their customers' expectations.

The study investigates the following two research questions:

²The CRA and RED do not apply to Medical Devices. The Product Liability Directive does.

- *RQ1*: How are connected medical devices patched within their operational environment at HDOs? and
- *RQ2*: What kind of challenges do HDOs and medical device manufacturers encounter during this process and how are they mitigated?

The remainder of this dissertation is structured according to three peer-reviewed studies, as depicted in Table 1.1. Following these three empirical studies, Chapter 5 summarizes and reflects on key takeaways for the overarching research question of this dissertation, considers implications for governance, and provides an outlook on potential future work.

Table 1.1: Dissertation outline.

Chapter	Publication
Ch. 2	Kustosch, L.F. , Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2023). "Measuring Up to (Reasonable) Consumer Expectations: Providing an Empirical Basis for Holding IoT Manufacturers Legally Responsible". In <i>Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)</i> . [133]
Ch. 3	Kustosch, L.F. , Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2025). "Regulating Smart Device Support Periods: User Expectations and the European Cyber Resilience Act". In <i>Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)</i> . [134]
Ch. 4	Kustosch, L.F. , Gañán, C.H., van Eeten, M.J.G., & Parkin, S.E. (2025). "Patching Up: Stakeholder Experiences of Security Updates for Connected Medical Devices". In <i>Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)</i> . [132]

2

WHEN SECURITY AND PRIVACY FAIL: UNDERSTANDING CONSUMER EXPECTATIONS

With continued cases of security and privacy incidents with consumer Internet-of-Things (IoT) devices comes the need to identify which actors are in the best place to respond. Previous literature studied expectations of consumers regarding how security and privacy should be implemented and who should take on preventive efforts. But how do such normative consumer expectations differ from what is actually realistic, or reasonable to expect how security and privacy-related events will be handled? Using a vignette survey with 862 participants, we studied consumer expectations on how IoT manufacturers and users would and should respond when confronted with a potentially infected or privacy-invading IoT device. We find that expectations differ considerably between what is realistic and what is appropriate. Furthermore, security and privacy lead to different expectations around users' and manufacturers' actions, with a general diffusion of expectations on how to handle privacy-related events. We offer recommendations to IoT manufacturers and regulators on how to support users in addressing security and privacy issues.

This chapter has been published as: **Kustosch, L.F.**, Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2023). "Measuring Up to (Reasonable) Consumer Expectations: Providing an Empirical Basis for Holding IoT Manufacturers Legally Responsible". In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)*.

2.1. INTRODUCTION

There are a growing number of consumer Internet-of-Things (IoT) devices in daily use in homes, such as smart speakers, smart lighting, and other home appliances now being offered with network connectivity. Flaws have been exposed in consumer IoT devices after release and purchase, such as security vulnerabilities and misconfigurations [197] and undisclosed data collection flows [60, 206]. How published flaws are addressed by their manufacturers is inconsistent – ranging from no response to security updates to, in rare cases, product recalls (as for smart cars [97]).

Home users have varying ideas on who they want to take responsibility for securing the devices before they enter the consumer market [101]. In parallel, government-level policymakers in various countries have set standards for consumer IoT security and privacy [63, 76], in an effort to reduce the problems that devices come with ‘out of the box’.

Existing efforts in academia and policy focus on boosting the baseline of security and privacy for consumer IoT devices. Still, problems do arise, and home users attempt to mitigate them in their own way when this happens [27, 173, 199]. It is uncertain whether entities in the consumer ecosystem other than users are providing adequate paths toward resolving these problems, where this includes the responsibility of the IoT manufacturer to fix issues or even refund a purchase. What is also not well understood is what support home users have come to expect of others when they learn that something has gone wrong with the security or privacy of their device. This raises questions around whether they have the same expectations for IoT devices as for the more familiar categories of smartphones or personal computers.

It is critical to understand the presumptions users make as to who they can turn to, as it should be that they can go to the right person for the right help, and do so easily and with some confidence that it is a predictable process. Would they assume first to have to go to the point of purchase [185], ask a (supposedly) ‘tech-savvy’ friend [187], or stop using the device altogether [27]? At present, issuing a software update is the easiest path for manufacturers, but even this patching is patchy, and does not always remediate inherent defects [197].

We conducted an online survey with 862 participants to study their expectations about the handling of IoT security and privacy events for products that they might own. We did so by presenting systematically varied vignettes. We answer a series of research questions: **(RQ1)** What do consumers expect how manufacturers *will* respond to emerging privacy and security risks with IoT devices?; **(RQ2)** What do consumers expect how manufacturers *should* respond to emerging privacy and security risks with IoT devices?; **(RQ3)** Do expectations differ across product types and threat events?; and **(RQ4)** How do participants evaluate the user’s responsibility to handle emerging privacy and security risks with IoT devices?

In the legal domain, *reasonable expectations* are critical to determining when a product or service can be considered defective [259] and thus trigger liability and product conformity regulation. While there is prior research into consumer expectations around Internet of Things (IoT) and smart devices [162, 212, 220], it is centred around normative expectations – that is, the preferences of consumers for how things should *ideally* be and which actors should *ideally* be responsible [101]. This does not capture what can reasonably be expected once something goes wrong with devices already in the mar-

ket [101, 105, 123]. We examine reasonable expectations by what is reasonable to expect (likelihood expectations), relative to what is hoped for (normative expectations), where the latter have been explored regularly in existing literature. Our main contributions are:

- We provide empirical insights on an important but understudied topic: What are consumers' expectations when something 'goes wrong' with the security and privacy of IoT devices?
- We extend ongoing user research on IoT security and privacy by framing users' needs in terms of what they realistically *expect* from device manufacturers relative to what they *hope* for. We find consumer expectations diverge between these two types of expectations, between privacy and security risks, and across device types.
- Our results provide a new angle for consumer protection policymakers and IoT device manufacturers when considering users' expectations, and we frame recommendations for addressing user needs to meet their expectations.

2.2. BACKGROUND AND RELATED WORK

Here we frame existing research on home users' experiences with IoT security and privacy against legal processes involving reasonable expectations. These are then considered alongside the expectations then placed upon other actors in the market, such as manufacturers and retailers.

2.2.1. EXPECTATIONS OF IOT SECURITY AND PRIVACY

There has been considerable research on consumer expectations for IoT security and privacy. This can include the features users expect for security[220, 263] and privacy[15, 146], but also the security concerns they would want a solution for[44, 105, 262]. Existing work conceptualizes expectations as *normative expectations*[79, 117] – that is, what users' preferences are for how things *should* be to minimise the potential for security and privacy problems to reach those users.

Normative user expectations have been captured as indicators of many preferences relating to consumer IoT devices: purchasing decisions relative to data access preferences [71], intentions to use devices relative to utility and data sensitivity [220], and approachability of security and privacy protection solutions [123]. Normative preferences are embodied most clearly in research on the *contextual integrity* [17] of data, regarding individuals' privacy preferences around the appropriateness of data flows involving IoT devices [3, 6, 11, 162, 212].

Alongside normative expectations, realistic expectations have been examined, albeit in limited scope. Zhang et al. [264] studied users' likelihood expectations of internet-connected security cameras with facial recognition capabilities and found that scenarios involving facial recognition prompted higher discomfort and more surprise. Furthermore, Gabriele et al. [88] prompted fitness tracker users about how feasible and likely a range of different threat scenarios were, finding that participants indicated a general optimism bias by underestimating likelihood of negative outcomes.

Here we move beyond risk perceptions and focus on what users regard as being reasonable to expect from different actors to resolve security and privacy issues with IoT devices. To the best of our knowledge, Haney et al. [101] provide the only account so far

that relates to expectations about responsibilities for ensuring the highest security and privacy of IoT devices. Participants framed ‘ideal’ situations wherein IoT manufacturers would be duty-bound to uphold the security and privacy of their smart home devices; at the same time, participants were unsure if manufacturers were in reality willing or able to do so. It is this distance between what *should* be done as a preferred ideal, and what *can be expected* as reasonable, that we study here.

2.2.2. REASONABLE EXPECTATIONS IN LAW

Expectations of consumers of a given product play a role in the domain of product liability and conformity laws. A concept originating in the United States, consumer expectations can be taken into account in product liability cases, when a ‘consumer expectations test’ is an option for the plaintiff to prove that the design of a product is defective [61]. This is the case if the product “*failed to perform as safely as an ordinary consumer would expect when used in an intended or reasonably foreseeable manner*.” In product liability cases, the plaintiff must thus prove that the expectations of a reasonable consumer were breached by the manufacturer.

The ‘reasonable expectations’ of consumers are relevant in other legal frameworks. For instance, the European Product Liability Directive [242] requires manufacturers of products – including IoT devices – to ensure that products conform to specific requirements. A product is defective, or regarded as not conforming to requirements, “*when it does not provide the safety which a person is entitled to expect, taking all circumstances into account*.”, as is also applied in EU courts (e.g., [13]).

Regardless of jurisdiction, the decision-making of courts is complex and context-dependent. Different factors can be taken into account to determine if a product conforms with requirements, such as product marketing and presentation, the baseline of comparable products on the market, or pertinent regulations and standards (e.g., [76, 230, 242]). Among these considerations, we find expectations of ‘ordinary’ or ‘reasonable’ consumers [149]. How consumer expectations and ‘reasonableness’ are conceptualized and determined lies ultimately with the court. However, case law commonly shows how expectations of ‘ordinary’ or ‘average’ consumers are considered in the verdict (e.g., [84, 177]), and regulations highlight the importance of the expectations of a ‘public at large’ [242].

To the best of our knowledge, courts have not drawn on survey evidence to inform their assessments of consumer expectations in the domain of security and privacy. That being said, we hope to provide empirical support for those assessments with our study. Our work connects to these legal notions via quantitative data on a large sample, to capture what consumers expect as opposed to what they consider desirable.

2.2.3. THE ROLE OF OTHER STAKEHOLDERS

Governments and regulatory bodies are at work to establish a basic level of privacy and security for IoT devices, such as consumer data protection laws like the GDPR [231] or CCPA [178], or regulations aiming at securing connected devices specifically such as the EU radio equipment directive [230]. There are several industry and government organizations publishing guidelines and voluntary standards to help manufacturers implement improved security and privacy into their products, e.g., from NIST and ENISA [63,

74, 76].

Looking at market actors, there are a few instances of product recalls, in case of serious security and privacy risks. These include smartwatch encryption for younger users [175], network-vulnerable smart security cameras [122], and vulnerable automobile software [97]. The most common response is the release of a software patch [8, 182]. Retailers serve a role as a contact point when a purchased device has problems or must be replaced [185], and this role is being recognised in some EU countries, e.g., the Netherlands [165]. Other emerging initiatives involve IoT product ‘labels’ [70, 73, 160] to guide consumers to purchase more secure IoT devices with more transparent details of the security and privacy features; consumer guides complement this, e.g., Mozilla’s ‘*Privacy Not Included’ [161]. We revisit how the manufacturer response informs the role of other stakeholders (such as retailers and governments) in section 2.5 based on our survey results.

Outside of policy and market mechanisms, Internet Service Providers (ISP) also well-positioned to detect, inform, or quarantine infected users [27, 42]. Otherwise, if a user has problems with a device, they may reach out to someone they regard as ‘informal’ technical support [187], or seek information on news or specialist websites [193] or on forums [111].

2.3. METHOD

To address our research questions defined in section 2.1, we deployed a vignette-driven online survey with 862 participants from the Prolific [190] crowdsourcing platform during August 2022. Participants were each presented with seven fictional text-based scenarios (*vignettes*) about a user experiencing a privacy or security risk with their IoT device, accompanied with varying ways in which the manufacturer and user respond to the situation.

As we were interested to measure expectations about concrete actions from manufacturers and users and how this would be influenced by IoT devices and security and privacy events, we systematically varied all these factors as variables in the vignettes to measure their relative impact on participants’ expectations.

Vignettes have been widely used in privacy and security user research [6, 11, 14, 19, 23, 148, 162] and allow to study participants’ judgments on multidimensional phenomena, while reducing social desirability biases of direct survey questions [130, 226]. Vignettes have been demonstrated elsewhere to be useful in exploring scenarios with *unresolved issues* in security and privacy [23], here being the uncertainty around who is best-placed to address a perceived security or privacy shortcoming with a consumer IoT device.

Each vignette described a situation with the same overall structure: 1) A user has an IoT device which gathers specific data and is used in a certain context and for certain purposes; 2) The user learns that the device has a software vulnerability (*security*) or is used for previously-undisclosed data practices (*privacy*); 3) The device manufacturer has responded in a particular way; 4) the user follows a certain course of action in response to the event. Each of these four phases constituted a factor in the vignette that could take on several varying levels, which are summarized in Table 2.1. An example vignette about a security event is given as follows, involving a protagonist (Alex); numbers in brackets

are inserted here (and do not appear in the survey itself), representing (1) Device, (2) Event, (3) Manufacturer response, and (4) User response:

Alex has several [1] **internet connected security cameras** at home, which are kept switched on continuously. The cameras continually collect video recordings of Alex's home and its surroundings to act as a deterrent against break-ins and allow Alex to check the video feeds remotely from a mobile app via an internet connection. Alex reads in a news post that a software vulnerability has been found in this device model and that similar vulnerabilities have been attacked. The [2] **vulnerability could allow other people to remotely install software on the device without Alex noticing**. The device could then be used to remotely attack other websites or devices connected to the internet, but Alex would still be able to use the device without noticing a problem. In response to this, the [3] **device manufacturer releases a statement** on their website and social media channels, which informs users about the vulnerability and the risks. Alex decides to try to [4] **return the devices** to the store where they were bought, hoping to receive a full refund or a replacement.

2.3.1. MEASURING CONSUMER EXPECTATIONS: THE IDEAL AND THE LEARNED

To measure participants' expectations about manufacturers' responsibilities and users' roles, we asked several 7-point Likert scale questions after each vignette, as follows:

- 1) **Likelihood expectation.** How likely a real manufacturer would respond this way;
- 2) **Normative expectation.** How appropriate the manufacturer response was. This relates to prior examination of what consumers expect of other ecosystem actors [101];
- 3) **Appropriateness of user action.** How suitable the user's action was in light of the scenario and manufacturer response;
- 4) **Vignette realism.** How realistic participants deemed the vignette to be.

This approach allowed us to simultaneously measure the impact of the vignette factors on these response scales. We designed two separate sets of vignettes, one for security events and one for privacy events, allowing us to contrast the arguably more state-driven nature of security dilemmas (whether a device is secure or not) with the context-driven nature of privacy dilemmas (whether personal privacy preferences have been respected).

2.3.2. SURVEY PROCEDURE

Participants on the Prolific platform were directed to a Qualtrics [191] survey, hosted at our research institution. After reading and agreeing to the informed consent, participants were presented with a short summary of consumer IoT devices to ensure all participants had a working understanding of what was and was not regarded as an IoT device (which is important for the purpose of shared understanding between researcher and participant [130]). To capture prior experience with internet-connected devices, participants were then asked to select from a multiple choice list of devices they have used at least once during the last four weeks.

Participants were each assigned a set of vignettes generated from source factors as in Table 2.1, constructed to resemble a scenario as in the example vignette (subsection 2.3.1). Participants then answered questions about a differing set of these kinds

Factor	Levels security vignettes	Levels privacy vignettes
Device	<ol style="list-style-type: none"> 1. Smart speaker 2. Smart watch 3. Smart washing machine 4. Smart security camera 5. Smartphone 6. Connected car 	<ol style="list-style-type: none"> 1. Smart speaker 2. Smart watch 3. Smart washing machine 4. Smart security camera 5. Smartphone 6. Connected car
Event	<ol style="list-style-type: none"> 1. DDoS 2. Unauthorized data access 3. Ransomware 	<ol style="list-style-type: none"> 1. Data collection without consent 2. Third party data sharing 3. Forced data collection
Manufacturer response	<ol style="list-style-type: none"> 1. Announce patch 2. Inform users via website and social media 3. No response 4. Recall 	<ol style="list-style-type: none"> 1. Announce update with more privacy settings 2. Inform users via updated privacy policy 3. No response -
User response	<ol style="list-style-type: none"> 1. Attempt to return device 2. Attempt technical mitigation 3. Seek advice online 4. Turn device off 5. Keep using as before 	<ol style="list-style-type: none"> 1. Attempt to return device 2. Attempt technical mitigation 3. Seek advice online 4. Turn device off 5. Keep using as before

Table 2.1: Overview of vignette factors and levels.

of vignettes. Participants either received a full set of security vignettes, or of privacy vignettes. Vignette construction is detailed in subsection 2.3.3.

After reading and answering questions about all assigned vignettes, participants were asked how confident they felt about their answers. Participants were then asked if the vignettes reminded them of any personal experiences with electronic devices, allowing them to provide personal stories [194] of security and privacy in an open text field. Participants then answered closing demographic questions, were debriefed and thanked for participation. It took 17.69 minutes on average to complete the survey (SD = 9.47 min), which also includes two attention checks; each vignette set included one Likert-scale question, which asked participants to answer with ‘agree’. After finishing all vignettes, participants were also asked to select a specific device from a short list of devices. The full survey instrument can be found in Appendix A.1.

2.3.3. VIGNETTE DESIGN

Participants were randomly assigned to either see security or privacy vignettes and were presented with seven vignettes in a random order to avoid sequencing effects [12, 203]. We opted to present seven vignettes to strike a balance between more repeated measures per participant (increasing statistical power) [12] while not mentally overloading them with too many vignettes [139]. After each vignette the four Likert scale questions described in subsection 2.3.1 were asked.

If participants rated the manufacturer response as inappropriate, the user response

as not suitable, or the vignette as unrealistic (selecting a value below the mid-value ‘*neither agree nor disagree*’), a free-text entry box was presented prompting to explain what motivated their answer. This encouraged participants to suggest other user or manufacturer responses that were not covered by the vignettes. These were typically seen to involve suggesting one of the response types presented in the survey, so for brevity these are not discussed further here.

Combining all possible combinations of the vignette levels depicted in Table 2.1 led to a total vignette population of 360 different combinations for security, and 270 different combinations for privacy. Adhering to methodological literature [12, 218], we reduced both vignette populations (*full factorial*) to a subset (*fractional factorial*), so that only a selected fraction of possible vignette factor combinations would be tested by participants. The resulting two subsets consisted of 91 different vignettes each and were then split up into 13 smaller subsets (*blocks*) with seven vignettes each, so that participants would be randomly allocated to one to the blocks. This kept the required number of participants manageable and limited the number of vignettes presented to participants to avoid mental fatigue [139].

We removed illogical combinations between factor levels to retain vignette credibility, e.g., the recursive example of a manufacturer updating the privacy policy to inform users about an updated privacy policy explaining additional data collection. We furthermore ensured that every participant would see each factor level at least once when reading the seven vignettes (e.g., not be predominantly presented with vignettes about smart cars, but see each device at least once) and that combinations of factor levels would occur equally often over the entire sample, e.g., to avoid that a recall of a smart speaker would occur more often than a recall of a smartphone.

We took great care in generating empirically grounded and realistic vignettes by deriving them from news reports, prior empirical literature, consultations with security and privacy as well as legal scholars, and a focus group. In the following paragraphs we explain our procedure and motivation for selecting the vignette factors and levels.

2.3.4. CHOICE OF VIGNETTE FACTORS

Choice of IoT devices. As factor levels, devices were selected which ranged from common ‘smart home’ devices such as smart speakers or IP cameras, to connected cars and smart washing machines. We also added smartphones as a prevalent and familiar device for comparison. The goal was to compare a diverse variety of IoT devices with varying usage contexts, data collection capabilities, and risks, to determine their influence over security and privacy expectations.

Choice of security and privacy events. We examined whether different security- and privacy-related risks would influence expectations on how manufacturers and users should handle them. We primarily based event types on prior user studies, and news reports. For instance, we identified reports of DDoS malware [153], unauthorized access to IoT sensor data [55, 201], and ransomware attacks targeting IoT devices [184]. Privacy-related events included reports of staff listening to device recordings for training of algorithms [60, 166], or device data being shared without the user’s consent [59, 206].

Events followed one of three different outcomes: that continued use of a device is im-

paired or ‘forces’ consent to be given; personal data from the device could be accessed by unknown parties (attackers or secondary data recipients), or; the device or its functional data is leveraged without the user’s knowledge or consent.

Choice of manufacturer responses. There is a focus in the literature on provision of software updates as a core response to security issues such as vulnerabilities [8, 182] and privacy issues such as providing more privacy controls [34]. We included these as possible manufacturer responses, but examination of news reports indicated a range of different responses beyond this. For instance, we noted product recalls in case of risks posed to children by smartwatches [175], smart security cameras being vulnerable to DDoS malware [122], or smart vehicle vulnerabilities [97]. There were also accounts of manufacturers not visibly responding directly to an event [48, 176, 183], reflecting that there is – as yet – little in the way of direct and consistent legal obligation for manufacturers to respond in a specific, predictable way.

Based on these reports and related research we conceptualized companies’ responses to disruptive events along a continuum, of enacting no responsibility to considerable responsibility [35, 36, 188], specifically: No reaction, informing users, releasing a software update, and recalling a device.

Choice of user responses. We grounded user responses in privacy and security user studies. However, empirical research on how IoT owners respond to security and privacy events is scarce [27, 199], as existing work mostly focuses on preventative mitigation by users [5, 87, 101, 105, 219]. We included five different user responses: 1) Keep using the device, due to e.g., discounting of risks to data [123, 136, 219, 262] or security [101, 219], or resignation [101, 136, 210]; 2) Unplug the device, ceasing or pausing use [27, 201, 224]; 3) Opportunistically seek help from others [57, 187] or online [10, 111, 201, 211]; 4) Attempt technical remediation oneself through device configuration or isolation from the network [87, 105]; 5) Request a refund or a replacement device from the seller. Such a response is commonplace when users perceive a defect in purchased goods and is protected by legal frameworks. However, with suspected security and privacy flaws this may be subject to the seller’s judgement and hence unpredictable.

UNCERTAINTY AND CONSUMER EXPECTATIONS

We phrased the vignettes so that the protagonist, and in turn the survey participant, would have incomplete information about the situation involving a security or privacy risk. For instance, all software vulnerabilities were phrased in a way that the vulnerability *could* allow for an undesirable outcome, or that data collected and shared with third parties *could* be linked to other information about the user. This level of ambiguity was chosen since users of consumer IoT devices usually face such uncertainty [27, 201, 224].

PILOT STUDY

Prototype vignettes were tested ‘offline’ in an iterative manner with volunteers without a technical background to check comprehensibility. This resulted in removal of illogical vignette combinations, language improvements, and efforts to give the protagonist a gender-neutral name (Alex).

A pilot study was conducted online with 32 participants from Prolific [190] to assess survey functioning, completion time, and vignette comprehension. It took participants 19.2 minutes on average to complete the survey, vignettes were rated as easy to understand, and open text responses did not indicate any major comprehension or technical issues. This resulted in slight adjustments to phrasing of some factor-level combinations within vignettes.

2.3.5. ETHICS

The study was approved by the host institution's human research ethics committee prior to survey deployment. To participate in the survey, individuals were informed that participation was voluntary, could be stopped at any time, and that no personally identifiable data would be collected. Participants had to agree to these points to be able to take the survey. We paid participants £3.00 for 20 minutes of their time, matching the minimum wage in the host institution's country.

2.3.6. PARTICIPANTS

Participants were recruited via the crowdsourcing platform Prolific [190] during August 2022. We screened for fluency in English, prior participation in at least five other studies on the platform, and a minimal approval rating of 95%. We did not screen for IoT device ownership or usage, but we did assess their experience, as we were interested if people with less or no IoT experience had differing expectations. In an effort to sample participants from different countries, we opened the survey several times, at different times and for different regions.

Demographics are summarised in Table 2.2. 862 participants took part in the survey: 443 female (51%), 399 male (46%), and 20 non-binary or no answer (3%). Age was skewed towards a younger population, which is a typical characteristic of Prolific samples [223]. Participants indicated to be from 30 different countries, which we mapped to regions for further analysis. The majority of participants lived in western countries (Europe and North America), while a smaller number lived in other regions such as Africa and Central and South America. Participants used on average 5.65 (SD = 2.30) internet-connected devices during the previous four weeks, indicating considerable experience with IT devices.

Due to random allocation to either the security or privacy vignette condition, participant characteristics (age, gender, region of residence, and device usage) were similarly distributed in both conditions. 23 participants got one of the two attention-check questions wrong; no participant failed both. We found no indication of suspicious response patterns from these 23 participants, and thus treated their responses as genuine and included them in analysis.

2.3.7. VIGNETTE AND RESPONSE QUALITY

Prior security-related studies have indicated the usefulness of realism checks for scenarios, for moderating the quality of response data [20]. We checked the responses to the prompt *'The situation described in the story is realistic.'* on a 7-point Likert scale, where a 1 would indicate *'Strongly disagree'* and 7 *'Strongly agree'*. On average, vignettes were rated to be realistic, not warranting concerns about implausible vignettes: for security

Age		Region of residence	
18-24	312	Europe (inc. UK)	464
25-34	303	North America	304
35-44	142	Africa (South Africa)	65
45-54	55	Cent. and S. America	26
55-64	42	Other	3
65	6		
Prefer not to say	2		
Sample size:			862

Table 2.2: Distribution of age and region in the sample. The three most prevalent countries were USA (N = 179), Canada (N = 125), and Portugal (N = 86).

vignettes, mean realism rating = 5.48, SD = 1.16; for privacy vignettes, mean realism rating = 5.67, SD = 1.09.

Participants' confidence in their responses was checked with 'How confident do you feel about your answers to the previous stories?', on a 4-point Likert scale from 1 = *Very unconfident* to 4 = *Very confident*. Participants were highly confident about their responses (Mean = 3.51, SD = 0.54).

2.3.8. DATA ANALYSIS

To answer our research questions, we first assessed average response patterns across vignette levels to identify general trends in the data. To quantify vignette factors' effect on expectation ratings, we ran multilevel regression models with maximum-likelihood estimation. Vignette factors were used as explanatory categorical variables predicting the response variables *Appropriateness of manufacturer response*, *likelihood of manufacturer response*, and *suitableness of user response*. Thus, six regression models were run, one for privacy and one for security for each response variable. In each model, we tested if demographic background (age, gender, region) and recent device usage had an effect.

Multilevel regression analysis allowed us to conduct tests of significance of factor levels, assess model fit, and control for any effects of participant characteristics such as recent device usage or region of residence. As suggested by methodological literature [12], random intercepts were included to account for individual differences between participants. As the response variables were on seven-point Likert scales, we treated them as continuous [172].

All regression models were built up with the following sequence: 1) A baseline with vignette levels as fixed explanatory variables and a random intercept term. For all tested regression models, likelihood ratio tests of the random intercept term were statistically significant, indicating that accounting for differences between participants explained significant variance in the data; 2) After the baseline model was defined, participant-level variables (age, gender, region, recent devices usage) and possible interaction terms were added in a step-wise fashion to assess whether they significantly improved model

fit. In Table A.1 in the Appendix we include the final models, reporting participant-level or interaction effects in the next sections only if they were found to be present.

Open-text responses were reviewed for any additional insight into participants' motivations behind their survey answers. We include representative quotes alongside results in the next section. To study participants' personal experiences with security and privacy incidents, two researchers independently reviewed the text responses to the survey question ('*Did the previous stories remind you of any personal experiences you have had with electronic devices?*'). During this thematic analysis [32], initial codes of reoccurring themes in the data were generated, which were then regularly discussed between the researchers in an iterative coding process.

2.4. RESULTS

2.4.1. EXPECTATIONS OF MANUFACTURERS

We first present how participants judged the manufacturer responses described in the vignettes, as an expression of expectations about how IoT manufacturers *would* and *should* respond to security and privacy events.

LIKELIHOOD JUDGEMENTS OF MANUFACTURER RESPONSES

Our first research question (RQ1) examines what consumers expect of how device manufacturers actually *will* respond to emerging privacy and security risks with IoT devices, as a construct closely relating to reasonable expectations. The left-hand side of Figure 2.1 shows how likely the manufacturer responses to a security event were rated on average across device types; Figure 2.2 does the same for privacy events.




Figure 2.1: Ratings of likelihood and appropriateness of manufacturer responses to security events. Measured on a 7-point Likert scale with 1 = 'Extremely unlikely' to 7 = 'Extremely likely' for likelihood of manufacturer response and 1 = 'Strongly disagree' to 7 = 'Strongly agree' for appropriateness of manufacturer response.

For security vignettes, patching was seen as the most likely response overall (mean of block 'Announce patch' = 5.70), followed by informing users about the risks (mean of

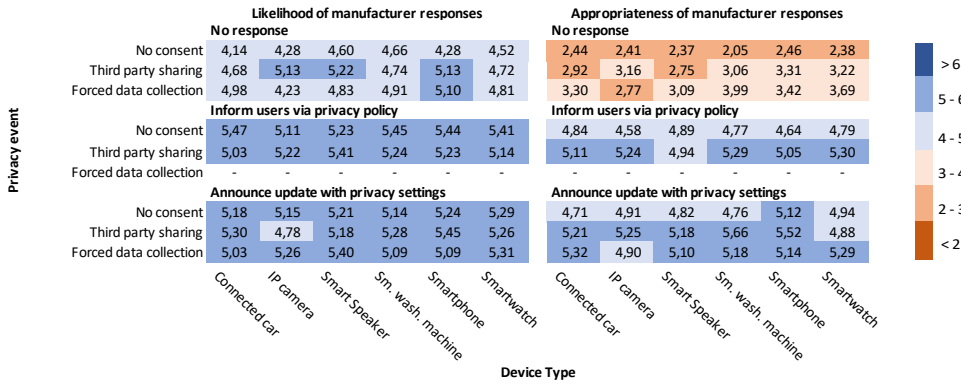


Figure 2.2: Ratings of likelihood and appropriateness of manufacturer responses to privacy events. Measured on a 7-point Likert scale with 1 = ‘Extremely unlikely’ to 7 = ‘Extremely likely’ for likelihood of manufacturer response and 1 = ‘Strongly disagree’ to 7 = ‘Strongly agree’ for appropriateness of manufacturer response. Empty cells correspond to less plausible vignette level combinations, which were removed from the design.

block ‘Inform users’ = 4.87), recalling devices (mean of block ‘Announce recall’ = 4.84), and lastly, not visibly/publicly responding at all (mean of block ‘No response’ = 4.05). Figure 2.1 illustrates this, as average ratings were generally higher for patching across security events and IoT devices. The unexpected nature of manufacturers not responding was reflected by participants’ comments, e.g., “I believe most manufactures would speak about the matter and possibly would recall the devices or issue an update for the devices.” (PID293). The regression analyses (Table A.1, Model 3) supported this trend: for security, all manufacturer responses were judged as significantly more likely than no response.

Figure 2.2 shows that for privacy vignettes, participants rated it most likely that a manufacturer would update the privacy policy and inform users (mean of block ‘Inform users via privacy policy’ = 5.28), while no response was seen as least likely (mean of block ‘No response’ = 4.74). In contrast to security vignettes however, this response omission was seen as relatively more likely. The regression analysis (Table A.1, Model 4) shows that the likelihoods of the two explicit manufacturer responses were comparable, indicated by similar coefficient estimates.

NORMATIVE JUDGMENTS OF MANUFACTURER RESPONSES

Our second research question (RQ2) examines what consumers expect of how IoT manufacturers should respond to emerging privacy and security risks with IoT devices to contrast such normative preferences with perceptions of the status quo. This is then closer to the aims of prior research [101]. The right-hand side of both Figure 2.1 and Figure 2.2 present how appropriate the manufacturer responses were, rated on average across both device types and security or privacy events respectively; Model 1 and 2 in Table A.1 in the Appendix show the regression models predicting appropriateness ratings of manufacturer responses.

For security vignettes, participants rated a product recall as the most appropriate manufacturer response across IoT devices and security events (mean of block ‘Announce recall’ = 5.76), followed by a patch (mean of block ‘Announce patch’ = 5.34). The manu-

facturer omitting a response to a security risk was rated as highly inappropriate on average (mean of block ‘*No response*’ = 2.05) across devices and security events, which was significantly lower than all other responses, as indicated by the regression coefficients in Table A.1 (*Model 1*). Among those participants who provided low ratings, indicative reasoning included, “*They are completely ignoring an issue that could put people in danger, if malicious people were to find out their location, for example.*” (PID283).

Both recall and patching received comparable ratings across device classes and security threats, demonstrating that participants valued both responses regardless of context. Several participants also stressed the importance of the timing of patches, e.g., “*An expected date of update would be appropriate, as well as some sense of urgency*” (PID169).

For privacy events (Figure 2.2), releasing a software update with more privacy controls was most preferred across devices and privacy events (mean of block ‘*Announce update with privacy settings*’ = 5.09). The regression model predicting the appropriateness ratings for privacy (Table A.1, *Model 2*) supported this, as both of the explicit manufacturer responses were rated significantly more appropriate than no response.

DEPENDENCY ON DEVICE TYPE AND RISK EVENT

In **RQ3**, we asked if participant expectations would vary across different device types and security and privacy risks. Looking at Figure 2.1 and Figure 2.2, this involved distinguishing between cells across rows and columns. For expectations of how manufacturers would actually respond, we did not find substantial effects of the type of security event on participants’ estimations. For privacy events however, we observed that it was rated least likely that a manufacturer would not respond after it became public that data was shared without consent (*no consent*), while it was seen as comparably more likely that a manufacturer would show no response after it became public that data is shared with third parties (see Figure 2.2).

For device types, recalls were judged most likely as a response for vulnerable connected cars, presumably since recalls of cars occur more often than for the other devices. It was rated least likely that a manufacturer would not respond to arising security vulnerabilities of security cameras and smartphones. *Model 3* in Table A.1 indicates that a product recall and a patch for a vulnerable smart washing machine were seen as significantly less likely than for a smart security camera. This could be driven by smart washing machines being seen as less critical or complex, with participants judging both recalls and patches as excessive and thus unlikely, e.g., “*It’s extremely rare that companies would make such expensive moves. These are absolutely the right things to do, but [...] it’s much more convenient to warn costumers, issue patches, or even ignore the problems.*” (PID140).

Normative expectations were also slightly influenced by IoT device type and the nature of arising security and privacy risks. No manufacturer response was rated as especially problematic for security cameras and smart speakers (e.g., “[*Alex*] should completely shut it [*Smart speaker*] down and wait until it is clear that the software patch is ready” (PID70)), and relatively less problematic for smart washing machines (e.g., “*Someone knowing my washing schedule really wouldn’t concern me. I’d probably just keep using it.*” (PID604)). Informing users of connected cars and smartphones was rated as less appropriate than for the other devices. We assume this was due to cars and smart-

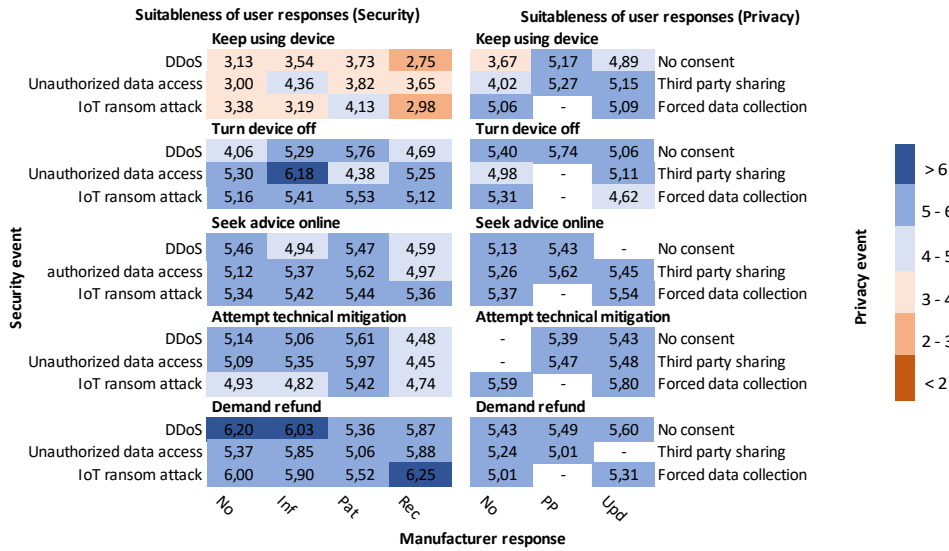


Figure 2.3: Ratings of user response suitability across security and privacy events and previous manufacturer responses. Measured on a 7-point Likert scale with 1 = 'Strongly disagree' and 7 = 'Strongly agree'. Empty cells correspond to less plausible combinations, which were removed from the design or combinations that were removed during the generation of the fractional factorial design, as not all three-way combinations of vignettes could be included. Key: No = 'No response'; Inf = 'Informed users'; Pat = 'Announced patch'; Rec = 'Announced recall'; PP = 'Informed via privacy policy'; Upd = 'Announced privacy settings'.

phones usually being needed on a daily basis, and only informing users was seen as not sufficient, e.g., "This solution [inform owners] does not seem proactive enough" (PID608).

Privacy risks also influenced the judgement of manufacturer responses. Scenarios where the user finds out that data has been collected from the device without consent ('No consent', Figure 2.2) negatively impacted how appropriate manufacturer responses were rated, especially for 'No response', as also shown in the regression model (Table A.1, Model 2, significant difference between no consent and third-party sharing).

2.4.2. THE USER'S ROLE

In this section we answer **RQ4** and report on how participants evaluated the behaviors exhibited by the user in the vignettes. These results do not only inform how participants judged the user's responses specifically, but also how these judgements translate to their expectations about the suitability of the user's options in reaching a satisfactory response to particular events. Figure 2.3 depicts how user behaviors presented in the vignettes were rated across previous manufacturer responses and security and privacy events and Model 5 and 6 in Table A.1 present the results of the regression models predicting the suitability of the user responses with vignette factors.

HANDLING SECURITY RISKS

On average, participants in the security vignette condition rated returning the product for a replacement or refund as the most suitable user action (mean of block 'Demand re-

fund' = 5.79), and continued usage as the least suitable (mean of block '*Keep using device*' = 3.46). When asked for alternatives for the user after giving a low response, explanations included, e.g., "*Simply turning off the device and ceasing to use it is a waste of money. Instead, Alex should return the smart speaker.*" (PID659). All user responses were rated significantly higher than *Keep using*, as indicated by the regression coefficients in Table A.1 (*Model 5*). Attempts by the protagonist to find a technical solution themselves were rated as less suitable (mean of block '*Attempt technical mitigation*' = 5.01) than simply turning the device off (mean of block '*Turn device off*' = 5.18). This apparent scepticism towards the user attempting a technical strategy was also reflected in text responses, e.g., "*Doing the configuration on his own requires specific knowledge and from this story I get the feeling that he doesn't have it himself. He should contact specialists and take time to decide what's best.*" (PID655).

Returning the device for a refund was seen as especially suitable in case of DDoS vulnerabilities (Figure 2.3), while keeping a device in this case was rated very low. This effect was also reflected in a significant regression coefficient of IoT ransomware in comparison to the reference DDoS (Table A.1, *Model 5*). It was rated highly suitable for the user to return the device when the manufacturer announced a recall. However, all other user responses, especially *attempt technical mitigation* and *keep using*, were rated lower if the manufacturer previously announced a recall. If a vulnerability allowed unauthorized access to sensor data and the manufacturer informed users about this, participants deemed it especially appropriate for the user to stop using the device (mean = 6.18).

HANDLING PRIVACY RISKS

User responses to privacy events were rated similarly on average, with a user attempting a technical solution as the most suitable (mean of block '*Attempt technical mitigation*' = 5.53), and the user continuing to use the device as the least suitable (mean of block '*Keep using device*' = 4.74). In comparison to security risks, continued use was rated much higher (mean of '*Keep using device*' for security = 3.46), indicating that keeping the device on after a suspected privacy-violating event was seen as a comparably more acceptable option than after an emerging security risk.

There were lower regression coefficient estimates for the user responses to privacy events (*Model 6*) than for security events (*Model 5*) (see Table A.1): in the case of privacy issues with IoT devices, participants were much less decided on a proper user response, to the extent that turning the device off was rated as the second most suitable response. This lack of a clear preference was also illustrated by a comparably low model fit (*Model 6*: $R^2_{\text{privacy}} = 0.172$ vs. *Model 5*: $R^2_{\text{security}} = 0.306$). Participants' comments hinted here at privacy resignation and feelings of helplessness, e.g., "*[The] decision isn't ideal but what alternatives are there? Alex could use an older-model "dumb" phone or look into a more security conscious manufacturer for a new device.*" (PID511).

The nature of the privacy event only slightly influenced responses: the user keeping the device was rated especially low if data had already been harvested without consent and the manufacturer did not respond (mean = 3.67). However, if the manufacturer informed users about the same privacy violation via an updated privacy policy, continued use was seen as much more suitable (mean = 5.17). This finding corresponds with the low appropriateness ratings participants gave all manufacturer responses to this privacy

event (*No consent*). In fact, participants viewed it as the best option for the user to demand a refund or turn the device off in this case.

2.4.3. PERSONAL EXPERIENCES

To relate participants' expectations elicited by the vignettes with their personal experiences, we analyzed the text responses to the optional question '*Did the previous stories remind you of any personal experiences you have had with electronic devices?*'. In total, 310 participants provided answers with a wide range of topics.

Of 310 participants, 58 provided accounts of how they experienced privacy or security incidents with their devices and how they or the manufacturer responded. The most commonly mentioned response was to stop using the device in some way ($n = 17$), like PID540, who noted: "*I stopped using a certain smart watch after it was unclear what data was collected from the manufacturer and third parties.*". Other variations of this included interrupting usage until the situation was perceived to improve: "*I stopped using my [smart speaker] after [news about data collection] came out about it, until [the manufacturer] gave me better control over my data.*" (PID382). This illustrates how users relied on manufacturers to respond and their willingness to pause device use until they received explicit reassurance. However, replacing devices in case of no manufacturer response was also seen as an option, e.g., "*When there were issues with cameras, I simply shut mine down and removed them for a time then switched over to something else that was more secure.*" (PID673).

Other participants ($n = 14$) described changing device or privacy settings, for example: "*[I] have had manufacturers of devices I've used update their privacy policy, also their data collection practices. I've modified my privacy settings according to the updated policies.*" (PID64). A few participants ($n = 6$) mentioned technical approaches such as limiting network capabilities, home network separation, or factory resetting. There were also rare stories of successfully having a device refunded or a device recalled: "*I had bought a [...] phone which had a security vulnerability [...] I had to return the phone [...] at the request of the manufacturer*" (PID70). There was also mention of directly contacting the manufacturer for support: "*[...] I saw many reviews stating that the speaker sells data collected from the speaker [...] I contacted the manufacturer who assisted [...] with instructions on how to turn on privacy settings*" (PID362). Generally, these reactions to security or privacy threats validated the chosen user responses in our vignette design and correspond to previous findings [27, 136, 199].

Apart from responses to problems with a device, several other themes emerged; (1) *Device linkage*; 24 respondents wrote about their concerns about apparent connections between information provided during device use and seemingly unrelated online activities. For example: "*Just seeing targeted ads that are clearly from one devices usage communicated to a different device in the household.*" (PID419). (2) *Data uncertainty*; 22 respondents described a general uncertainty about privacy policies and data flows (e.g., "*I have several smart devices [...]. There are times I don't believe there is enough transparency about how this data is used, stored, or sold. I have felt companies are dishonest about these issues before which makes me hesitant to continue to use smart products sometimes.*" PID366). (3) *Dilemmas*; 20 respondents felt concerned and experienced dilemmas about whether security and privacy risks should be accepted, either in the form of resignation

(e.g., “[...] I am feeling helpless [about data collection], as there is nothing I can do about it, so I can either stop using the devices or use it and be ‘tracked’ down.” PID137) or as a convenience trade-off; “[...] companies sharing the information has crossed my mind. But at the end of the day, there’s not many ways around it, using the device is still more convenient than not using it.” (PID220). These themes reinforce the findings in Section 2.4.2, as they demonstrate a general uncertainty about data flows and how IoT users should manage privacy.

2.5. DISCUSSION

Here we revisit our research questions and situate our findings within prior literature and ongoing discussions.

What manufacturers *are likely to do*. In RQ1 we asked what consumers expect how manufacturers *would actually* respond to emerging IoT risks. As indicated by Figure 2.1, we found that participants in our study expected manufacturers to patch security vulnerabilities in IoT devices. This resonates with the current focus in policy circles. In contrast, no response at all was seen as unlikely, indicating that manufacturers are expected to visibly respond if a security event occurs. This supports recent standardization efforts recommending that IoT manufacturers notify and communicate with users in case of security incidents [76], and highlights the position of ISPs as being able to triangulate security problems to specific users (e.g., [27]).

The picture was less clear for emerging privacy issues with IoT devices (Figure 2.2), as different manufacturer responses were rated as comparably likely and no response was seen as somewhat less expected. A manufacturer not acting on problematic data flows was seen as highly inappropriate yet very conceivable. This hints at a lack of consumer trust despite GDPR regulations [178, 231] and a learned helplessness and resignation regarding control over the occurrence of privacy violations, and is in accordance with prior work [101, 136, 210].

These findings provide legal scholars and policymakers with novel empirical perspectives on the notion of consumer expectations in case of IoT security and privacy events. By using a shared language (*‘reasonable expectations’*), we show how it was expected by participants that manufacturers would patch security vulnerabilities or at least respond in some visible way. As discussed in subsection 2.2.2, liability case law is based on a case-by-case assessment, yet our findings can serve as a reference for the design of IoT security and privacy regulations (which do play a role in courts, see e.g., [259]) and provide new insights for legal scholars and practitioners on how the abstract notion of consumer expectations can be understood empirically.

What manufacturers *should ideally do*. Turning to RQ2, and how consumers prefer manufacturers *should* respond to emerging IoT events, we found that participants generally considered recalls and patching to be appropriate responses to security threats (Figure 2.1). Interestingly, patching was the only manufacturer response that was considered both appropriate *and* likely. Seeing patches as reasonable does rest on all security issues being resolvable by patches, without further manual fixes by the user, which in practice is often not the case [27, 199]. In contrast to patching, recalling was seen as

the more appropriate response, yet also considered relatively unlikely, even less likely than simply notifying users. This suggests a gap between consumer preferences and expectations.

As patching is much more prevalent than product recalls, consumer expectations appear aligned with, and perhaps habituated to, observed market behavior. This also fits with seeing a car recall as more likely than for other consumer IoT devices. Thus, expectations might change in the next few years, where stricter regulations could trigger more frequent recalls (e.g., not complying with minimum security requirements as in the upcoming delegated EU Radio-Equipment Directive will lead to the IoT device's removal from the market [239]).

For privacy, participants favoured it when a manufacturer announced the release of a software update with more privacy controls while also judging a lack of response as least desirable (Figure 2.2) (reinforcing prior findings elsewhere [260]). Notably, announcing a privacy software update and updating the privacy policy were rated as similarly appropriate. As with security events, this requires notification to be visible – in this case, within the device and/or companion app itself. Prior work has indicated that more control does not necessarily lead to higher trust in privacy [260], with a view to governments needing to enforce what manufacturers can and cannot do. This could also hint at a general loss of trust towards manufacturers to handle personal data appropriately, where more privacy controls would not help to restore the trust.

Our results build on previous work on IoT consumers' expectations of the responsibility of manufacturers and users, in which users expressed uncertainty if manufacturers would realistically meet their preferences [101]. Our results indicate that there are indeed discrepancies between consumers' preferences and predictions, as well as more clearly expressed expectations about security (manufacturers will likely patch and are unlikely to do nothing) than for privacy (with less clarity as to how manufacturers will likely respond). We furthermore broaden prior research on users' preferences on IoT security and privacy (e.g., [101, 105, 162, 220]) by contrasting normative preferences with a 'reality check' of expectations of actual likelihood.

Managing *different* security and privacy circumstances. Regarding RQ3, we found that different IoT device classes had an effect on what responsibilities participants expected from manufacturers. For instance, expectations around smart washing machines were less strict than for security cameras or smart speakers, which could be due to the device's less sensitive data. This matches prior work on privacy perceptions of IoT devices [73, 220]. For devices important for daily use (e.g., smartphones and connected cars), participants preferred a proactive response by the manufacturer beyond only informing them. Remarkably, connected cars did not cause a different effect. Compared to other device types, participants didn't see it as substantially less likely or less appropriate for a car manufacturer to not respond to security vulnerabilities, even though these can conceivably lead to safety hazards. For privacy events, manufacturer responses were rated as less appropriate for vignettes describing that data was harvested from the device without consent, which implies that this privacy violation reduced appraisals of manufacturer responses regardless of the actual response. Previous work has established the importance of user consent [6, 11, 44, 162, 211], and our results extend this notion by demonstrating

how the breach of this fundamental privacy principle also negatively affects subsequent efforts of the manufacturer to remediate.

2

How best to *involve users*. For RQ4, the user's involvement in addressing security and privacy risks was assessed (Figure 2.3). For emerging security risks with IoT devices, participants deemed it most preferable for the user to return the device for a refund or replacement. Depending on local legislation, the warranty period, and the seller's leniency, this might constitute a feasible path. However, as paralleled by several participants' comments, this route is arguably rarely observed in real life, and the chances of a successful return depend on many factors outside of the user's control. A recall notice would signal the feasibility of the response, but manufacturers might not have a reliable way of getting the notice to users.

Simply keeping a device in use after learning about a security problem was generally judged as highly ill-advised for the user. This perception was different for privacy, where it was seen as much more acceptable to keep the device on, especially if the manufacturer updated the privacy policy or announced an update with more privacy controls, despite the same prior privacy violation. This contrasts with prior research implying that users would turn off a device as if 'stopping a leak' [260] and illustrates how perceptions of privacy change with manufacturer signaling, but also as how limited the user's options were perceived.

It may be that IoT users are simply lacking options for action and control (as has been seen for both security [27] and privacy [105, 123]), making it a conceivable response for users to continue using the device, as unplugging could be undesirable due to discontinued operation, demanding a refund is seen as futile, and personal technical mitigation as unpredictable. That said, users' technical attempts to mitigate privacy risks were seen as more suitable than for security vulnerabilities.

These results also broaden prior findings of instances of users stopping use of their devices after (suspected) security risks [27, 201, 224], as we observed that turning IoT devices off was seen as a generally suitable response for both privacy and security risks, and was most frequently mentioned by participants as a previously applied response. That such a drastic step was seen as a suitable response illustrates how limited users' options appeared to be for a clear path to resolution, which highlights the necessity of actors better positioned to handle these risks to be involved.

If users were to stop using a device, this is difficult for those with expertise to detect, even if it at least stems some threats. This may also encourage a somewhat 'silent' departure from the smart device market (hinted at in Figure 2.4.2), where one 'bad actor' then tarnishes all reputations. This is arguably why consumer IoT devices are *generally* seen as lacking appropriate security (and requiring standards) although many devices exist which are already secure. Participants appeared just as amenable to stopping device use after a privacy issue as they were to demand a refund – this is then in the interests of manufacturers if they want to retain customers.

Prior work has also suggested that responsibility for protecting privacy of IoT devices was seen more with the manufacturer than with the user, while for security, the responsibility of the individual user was also central [101]. This could further explain why in our study, participants seemed to have clearer expectations of appropriate ways for the user

to handle security risks (try to get a refund, and avoid continue using a device) than for privacy, as the manufacturer is seen as responsible for remedying technical problems.

2.5.1. RECOMMENDATIONS

Here we list future directions and recommendations for ecosystem stakeholders.

Establish post-purchase maintenance and support. IoT users generally expect and appreciate explicit responses from manufacturers, preferably more than just a warning, which might remain unseen and be perceived as insufficient. Participants also voiced how they would switch brands or return devices in case manufacturer handling of security and privacy would lag behind their expectations. To establish user support and trust for the post-market phase, manufacturers should follow standards such as from NIST [76] and keep an active communication channel with their customers. While effective communication is not trivial to achieve, a collaboration with ISPs to reach identifiable customers could also be a fruitful direction.

Smooth the path for predictable outcomes. As governments are also seen to hold responsibility for IoT security and privacy [101], our findings furthermore provide regulators with insights into consumers' expectations. We recommend that regulators support users with routes for resolution that are coherent and predictable, such as specific and easily accessible advice. Furthermore, it is paramount to provide robust consumer protection laws to reduce incidents in the first place, but also to have regulatory or economic processes in place to incentivize appropriate and effective responses by device manufacturers, including smoothing the path for potential product returns.

Gather evidence with a view to its wider uses in law. In law, reasonable expectations are a fluid concept. There are no objective thresholds; the EU and US jurisdiction rely on the judge to interpret consumer expectations in each case. Our study offers concrete measurements of this construct to both legal practitioners and legal scholars in the product liability field, who might face questions surrounding consumer expectations of IoT devices in their work. A multi-disciplinary approach, in which empirical computer and social sciences support legal scholars with insights around assumptions about technology and its users, could constitute a promising future direction of academic work.

2.5.2. LIMITATIONS

While this study's sample is considerable in size, it is not representative of any specific national or global population. Due to Prolific's participant base, participants were mostly from 'western' countries. Furthermore, the sample was skewed towards younger cohorts, which is also typical of Prolific samples [223]. During sampling, we were interested in gathering a breadth of different regions and legislations and not in modeling any specific population. This limits the generalizability of the results yet nonetheless provides novel insights into consumer expectations across different regions.

The vignettes were bound to a limited number of factors, yet other aspects could also influence expectations. In all vignette permutations the user learns about security or privacy risks from a news post, while there are several other sources for users to learn about

possible security and privacy issues [73, 193], such as word-of-mouth, unusual device behaviour, or direct notifications (e.g., by ISPs [42]). We opted for the news post as this is a common channel for home users [58, 193] and may be communicated itself by word-of-mouth or analogy [194]. Furthermore, including the price of the IoT device could have influenced expectations, with cheaper devices perhaps being seen as more vulnerable and premium products leading to higher expectations. However, adding more contextual factors to the vignettes' factorial design would have led to an explosion of factor level combinations. Thus, we encourage future work to explore such directions.

Finally, participants had to judge a fictional user's actions, such that it needs to be determined if this judgement would translate into actual behaviour on their side, though text answers imply that participants had similar experiences to those captured in our vignettes, as presented in subsection 2.4.3.

2.6. CONCLUSION

Using a vignette survey with 862 participants, we found differing expectations around the responsibilities of users and manufacturers how arising security and privacy events would and should be handled. Future work should look at other factors related to product liability law however, such as the state of the market and behavior of competitors. Future work should also go beyond the vignette factors considered here, to explore the impact of other factors on expectations, e.g., duration of device ownership and price, warranty conditions, and timeliness of manufacturer response.

3

PREVENTING FAILURES: EXPECTATIONS OF SECURITY SUPPORT OVER DEVICE LIFETIMES

Supporting consumer IoT devices with updates is crucial to ensure their security. However, this support period is usually shorter than the device's actual lifespan, resulting in millions of unsupported and vulnerable devices. The upcoming European Cyber Resilience Act (CRA) addresses this by requiring manufacturers to support their products for the expected use time, which should be based on reasonable user expectations. In this work, we thus empirically explore the concept of user expectations regarding smart devices' use times and security provision by conducting a large-scale survey in five EU countries (n = 993). We find that respondents' smart device use times and lifetime expectations exceed the CRA's baseline of five years for a majority of device categories and vary substantially across device categories, their "smartness", and individuals. Respondents also consider different factors for the lifetimes of smart and conventional devices. Surprisingly, a majority of respondents expected update support to correspond with devices' full lifetimes, highlighting how the current market dynamics of short support times seem to contrast expectations. Our results provide novel insights for manufacturers and market authorities who will need to determine support periods for smart products in the coming years.

This chapter has been published as: **Kustosch, L.F.**, Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2025). "Regulating Smart Device Support Periods: User Expectations and the European Cyber Resilience Act". In *Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)*.

3.1. INTRODUCTION

To secure the growing population of consumer IoT devices from an evolving threat landscape, continuous support over the product's lifecycle is crucial. However, security updates for such smart devices often end prematurely and their presence or duration is often not disclosed to consumers at purchase. A recent report from the US Federal Trade Commission found that for 89% of the reviewed smart products, no disclosure on the duration of updates support was provided [81], seemingly a status-quo in the smart device market [189, 196].

Various regulatory and policy initiatives aim to increase transparency of the software support duration to provide certainty to consumers about the security capabilities of the devices that they purchase. Within this, there is an expectation that consumers will gravitate toward more secure devices within purchase decisions. The US Cyber Trust Mark announced in 2023 is a voluntary product labeling scheme where manufacturers of smart devices disclose the support period as part of the label [80, 152]. In the UK, the PSTI act *requires* manufacturers to disclose a minimum support period for security updates [229] to comply with the law, instead of the voluntary label approach in the US.

While *disclosing* information to consumers for a more informed purchase is an important step, the question then arises of *how long* this support should actually be provided. The approach of the European Commission is embodied in the Cyber Resilience Act (CRA)[243], launched in 2024 and coming into force in 2027; to the best of our knowledge, the CRA is the first regulation that explicitly sets legal product requirements on a *duration* for how long security measures must be provided across a broad range of products. Specifically, security vulnerabilities of “products with digital elements” must receive security updates for their *expected use time* but at least five years¹.

In the determination of the expected use time, that is, device lifetime, reasonable user expectations play a crucial role. Manufacturers will have to actively determine this period based on users' expectations and disclose how they considered this in the product's technical documentation. Similarly, market surveillance authorities will check such documentation and, in turn, determine what they see as a reasonable user expectation to determine compliance with the CRA. If they find that the support duration for a product violates expectations, they could remove the device from the EU market.

Thus, the user perspective, previously an afterthought in security support duration, suddenly takes center stage in the regulatory realm with global implications. Manufacturers sell the same smart devices in other markets than the EU and could streamline global compliance efforts and release the same security updates as in the EU, following the “Brussels Effect” [29, 142, 174]. However, the abstract concept of reasonable user or consumer expectations regarding a product's expected use time bears uncertainty for smart device manufacturers and market surveillance authorities. For instance, is it reasonable to expect that a smart washing machine will be used for longer than a smart-watch and should thus be patched for longer?

Prior research only provides limited insights into the use time or update support duration of smart devices. While [103] studied smart home users' perspectives on the end of support and [133] empirically measured reasonable consumer expectations regarding

¹With exceptions possible (Article 13.8. and Recital 60).

security and privacy incidents with consumer IoT devices, the expected or observed *duration* of device use or update support has not been directly measured, although [160, 164] find that longer support periods were generally preferred by consumers.

As the CRA explicitly links the duration for security support to users' expectations of the product's lifetime, we aim to fill this gap by conducting a large survey among EU consumers to empirically assess their expectations and behaviors regarding smart device use times. We ask the following Research Questions: (i) How do consumers use their smart devices, and for how long?; (ii) How long do consumers expect different smart device categories to last, and which factors influence these expectations?; (iii) How do consumers perceive security and software update support over smart devices' lifespans?; and; (iv) Are there differences among EU member states regarding consumers' smart device usage, expectations, and security and software support perceptions?

We make the following contributions:

- We conduct an online survey with 993 participants in five different EU member states and collect empirical data on and explain a methodological approach to a concept of high relevance in the coming years: Smart device use times and users' expectations regarding them. Thus, we provide an empirical basis for manufacturers and market surveillance authorities how to conceptualize and approach users' expectations regarding smart product lifetimes.
- We measure smart device *use times* by studying 2753 individual smart devices and measure lifetime *expectations* via smart device vignettes. While expectations varied across devices and individuals, we find that, in aggregation, almost all studied smart products were expected to last longer than the CRA's baseline of five years, despite respondents factoring in aspects like fast innovation and replacement, incompatibility, or planned obsolescence. Additionally, a majority of respondents expected update support to correspond to the device's full lifetime, highlighting how the CRA's provisions seem to meet expectations already.
- We extend previous work on smart devices' end of support by focusing on the concept of security support *duration*. We find that respondents were aware of security updates being crucial to mitigate security risks, although updates were generally opaque (non-perceptible), and their eventual end was not a strong trigger to stop device usage - better alternatives or declining functional performance were much stronger motivators.

In the remainder of the paper, we first consider the regulatory underpinnings and related work on users' expectations regarding smart devices, followed by our survey methodology and results. We then contextualize our findings in the wider regulatory and academic field and provide recommendations for manufactures and market authorities.

3.2. BACKGROUND AND RELATED WORK

Here we frame existing IoT users' experiences with IoT security and privacy against legal processes involving reasonable expectations. These are then considered alongside the expectations then placed upon other actors in the market, such as manufacturers and retailers.

3.2.1. EU SUPPORT DURATION REGULATIONS

User or “consumer” expectations are often part of legal considerations. In the context of cybersecurity, the most prominent piece of current legislation in the EU is the Cyber Resilience Act (CRA), adopted at the end of 2024. It regulates manufacturers that place “products with digital elements” (i.e., software and hardware products) on the EU market. It thus has global impact, as any manufacturer must adopt its measures when placing products on the EU market, regardless of their location. Internet of Things (IoT) devices are prominent examples of products with digital elements; they must adhere to numerous cybersecurity requirements (e.g., security-by-default, Annex I(2)(a)). The CRA reserves a special position for certain categories of IoT devices. For instance, smart home security devices and virtual assistants are an “important” category due to larger security risks than average products with digital elements.

The manufacturer must, as part of a comprehensive list of requirements, ensure that any detected vulnerabilities are swiftly mitigated for the entire “support period” of their device. Manufacturers must determine this support period themselves, based on several criteria provided in the CRA (Article 13(8)). The manufacturers *shall* (i.e., must) consider in this determination: “reasonable user expectations”, the nature of the product (including its intended purpose), and other EU legislation that may already determine lifetimes for the product category. Additionally, manufacturers *may* take into account support periods of similar products, the availability of operating environments, the support period for integral components of the product created by third-parties and guidance provided by administrative authorities (the “ADCO”). Manufacturers are thus obligated to take into account how long users would reasonably expect their product to remain in use, i.e., the product’s expected lifetime.

Manufacturers must include the chosen support period with the product’s documentation (Article 13(19)). Furthermore, they must share the information on which they base the chosen support period (Article 13(8)). Market authorities must monitor chosen support periods and take corrective action when they consider the support period incorrect.

European product legislation often connects product support periods and user expectations. For instance, the Sale of Goods Directive (SG Directive) offers another example of this approach [237], aiming to protect consumers when signing a sales contract. It requires that goods with digital elements (e.g., smart products) remain in conformity with the sales agreement. For instance, a video app on a smart TV cannot suddenly lower its picture quality, if the initial contract noted a higher quality. It thus requires the product to receive updates, including security updates, “for the period of time [...] that the consumer may reasonably expect” to conform with the contract (Article 7(3)). However, in contrast to the CRA, the SG Directive does not place a burden on the seller to provide how they assessed their support period and to share it with consumers and market authorities.

Thus, the support period determined by the manufacturer plays a key role for smart devices in the EU. Manufacturers and market authorities must know what support periods are reasonable for a breadth of product categories. Consumers, at the same time, can compare support periods so they can make informed purchase decisions. Given these dynamics introduced by the CRA, it is vital to understand user expectations regarding support periods of products with digital elements.

3.2.2. CONSUMER EXPECTATIONS AND PERSPECTIVES OF IOT SECURITY AND PRODUCT LIFESPANS

Prior work has studied consumer expectations of security support for IoT devices [133] – this included manufacturer liability and responses in case of particular incidents, finding that expectations differ per device and that, when measuring expectations in the legal context, it is important to distinguish between reasonable (realistic) and normative expectations (ideal). Here we study conformity to update support commitments, as in the CRA.

Concerning software support for IoT specifically, there is great uncertainty among users, even when directly presented with pertinent information (such as “security labels”) [102, 103], for instance that users say updates are important, but do not understand the implications of using unsupported devices or what end of support means [70], all the while lacking clarity as to the purpose of updates [104]. These outcomes point to users needing assurance in a market that is full of uncertainty – this suggests there could be value in not only setting a support duration, but strengthening assurances that it will remain over time, as a reliable “minimum” duration.

Further, when software support duration is considered as a product label attribute during device selection [73], prior work suggests there is a demand among consumers for the delivery of security updates for IoT devices for a longer period of time than they are accustomed to (e.g., [160, 164]). Thus, where previous work has suggested that support assurances would be valued, it lacks a clear connection to the expected duration of update support and, most importantly in the context of the CRA, the *expected* lifetime of IoT products, as [160, 164] measured *preferred* update support duration with predetermined values by the researchers (e.g., 2 years, 6 years, and lifetime).

However, consumers’ expectations of products’ lifetimes have been measured before, just not in the context of security and privacy, but in the domain of sustainability [54, 92, 93, 170, 181, 257]. For instance, [54, 92] and [181] measure expected lifetime or perceived longevity importance for product categories like mobile phones or vacuum cleaners. However, smart devices and the role of software updates and security were not directly studied. This research highlights different forms of product lifetime expectations, namely *intended* lifetime (how long a product is intended to be in use by its current owner), *ideal* lifetime (how long a product should ideally last), and *predicted* lifetime (how long a product will likely last). A conceptual piece by Bradley et al. examines similar concerns [30], outlining the conflicting interests between continued security patching and device longevity, proposing a paradigm shift: an architecture and differential responsibilities for different actors to keep devices running and secure for as long as possible. Here, we examine how user perceptions inform the provisions needed in determining support duration relative to the CRA.

In examining the electronics repair industry, Ceci et al. [40] found that 12% of their 112 survey participants stated that a smartphone or tablet never broke; however, the cost and hassle of getting a device repaired were major factors where participants had a device that developed problems but that they did not seek to get repaired. In studying views on IoT device obsolescence, Vats et al. [248] found that when interview participants were faced with external events such as an IoT device ceasing to function, needing an update, or losing smart features, they would try to find terms in the user agreement

that would force the manufacturer to take responsibility (and feel a lack of agency if they could not); this indicates an existing gap that CRA assurances could fill. The interview participants of Haney et al. [101] referred to an unspoken agreement at device purchase, that a manufacturer should protect buyer security and privacy, but that participants differed in how sure they were that a manufacturer could uphold those expectations, trusting larger firms as being more competent to do so.

3.3. METHOD

3

To explore user expectations regarding smart devices' lifetimes, we conducted an on-line survey with 993 participants recruited on the crowdsourcing platform Prolific [190] during November and December 2024. As the CRA is an EU legislation, we recruited an EU-based sample from France, Germany, Poland, Spain, and The Netherlands. We selected this subset as countries belong to the largest member states, together account for over half of the full EU population, vary in geographic location, culture, and markets, and to retain an appropriate sample size per country for robust statistical comparison. Due to the CRA requiring an EU-wide uniform approach, we were interested to study the potential variety of market dynamics across countries, such as user journeys, purchase channels, and experiences with smart devices and their security, potentially requiring differentiated support structures for users.

In the following, we detail our survey design, data collection and analysis, and the resulting sample.

3.3.1. SURVEY DESIGN

We first explain some of our conceptual considerations around the notions of consumer expectations and device lifetime, followed by survey design and preparation for publication in different countries.

CONCEPTUALIZING DEVICE LIFETIME EXPECTATIONS

We aimed to align with the CRA definition of product lifetime, which is “the length of time during which the product is expected to be in use”, which should be based on “reasonable user expectations, the nature of the product, including its intended purpose”. Previous empirical work on consumer expectations provided a solid methodological basis for phrasing this question in a survey. Gnanapragasam et al.[93] stress the importance of clearly differentiating between three forms of expected product lifetimes: *Intended* (how long a product is intended to be in use by its current owner), *ideal* (how long a product should ideally last), and *predicted* (how long a product will likely last). Other related work on smart devices' consumer expectations highlights the necessity to differentiate *normative* (ideal) and *reasonable* (realistic) expectations and argues that realistic expectations correspond more closely to the legal notion of reasonableness, rather than normative ideals [133].

Thus, we conceptualized user expectations about product lifetimes according to the following two criteria:

- The expectation is about a *predicted* period of time for how long a given smart device will remain functional to be used for its intended purpose, not what is desired

for how long it *ideally* should last. i.e., the expectation should be reasonable.

- This period should be independent of ownership, as devices can be sold or gifted onward [140], and we did not want to measure the initial use cycle, i.e., a user's preference when to move to a newer device.

Thus, when asking about the expected lifetime of a smart device, our question was: “*If you had to predict, for how many years do you expect such a device to last?*” (Q5.1.1). However, we also considered *intended* product lifetimes (how long a product is intended to be in use by its current owner) by inquiring with participants about how long they plan to continue using their *own* smart devices to then form a measure of their use time, adding up the durations of past use and intended future use.

MEASURING DEVICE LIFETIME EXPECTATIONS

To quantify expectations about smart devices' lifespans, we designed vignettes for several prototypical devices and presented survey participants with a random subset of them to measure how long they would expect the device to last. Table 3.1 depicts our selection of devices, and the vignettes can be found as part of the survey instrument in Appendix B.1.

Vignettes are short descriptions or scenarios, to provide respondents with context to elicit a response and are commonly used in security and privacy research [6, 11, 14, 133, 162]. Our vignettes described the device at hand, its high-level features, how it can be used, and its smart capabilities to ensure that respondents without any experience with it could also form a picture of this product. We followed common practice in security user research [103, 104, 133, 160], and selected a diverse range of smart devices with varying use cases, price points, and security and privacy implications and perceptions, which should all fall under the CRA.

As many white goods products become increasingly smart, we also tested whether expected lifetimes would differ between smart and conventional versions. We were interested in how participants would factor in the additional aspects introduced by the “smartness” into their expectations, such as software updates, online connectivity, and security or privacy risks.

We did not include a conventional version for all smart devices, as for some, we considered them an entirely different product based on their primary use case. For instance, we did not consider a regular watch (primary use: showing the time) a fair comparison to a smartwatch (primary use: running applications like health tracking, or messaging/notification).

SURVEY DESIGN PROCESS

We drafted the initial survey based on our research questions and the content examined for the literature review (Section 3.2). We also gathered feedback during a workshop with thirteen experts working in the fields of IoT security and law where the survey draft was presented, jointly discussed in the group, and subsequently improved regarding clarity, wording, and specificity of the IoT devices.

At this point, we involved policymakers and conducted four online meetings with policymakers from different EU member states, all involved in the drafting or future enforcement of the CRA. In the meetings, we inquired about their interpretations of

Device type	Product category	Conv. Comp.
Smart speaker	Entertainment	No
Smart thermostat	Home / Energy management	Yes
WiFi Solar inverter	Home / Energy management	Yes
Smart camera	Home security	No
Smart doorlock	Home security	No
Home router	Network equipment	No
Connected printer	Office equipment	Yes
Smart Smoke Detector	Sensors	Yes
Smart watch	Wearables	No
Smart washing machine	White Goods	Yes
Robot vacuum	White Goods	Yes

Table 3.1: Selection of devices used in the vignettes. “*Conv. Comp.*” denotes if a comparison to the conventional version of that product was made.

the CRA’s support period provisions and asked for their involvement in the study. We received feedback on our survey and suggestions for questions relevant to their work from 3 of the 4 organizations and then improved the survey regarding wording and flow and added a question about the purchase channels of respondents’ own devices (Q2.4 & Q2.5), as this was of interest for several policymakers.

A pilot run of the survey was conducted with 25 participants recruited on Prolific to check for any technical problems, question understanding, completion time, or any answer options not provided in our multiple-choice questions (i.e., Q4.3, Q4.4, Q7, Q9.1, Q9.2). Participants took 14.8 minutes on average and no technical issues emerged. Only minor adjustments were necessary for some questions for clarity, despite the pilot survey allowing respondents to indicate where they felt unsure about their answer.

After the pilot study, the survey was translated from English into the five countries’ respective languages (French, German, Polish, Spanish, and Dutch), by first being translated using Qualtrics’ native automatic translation tool based on Google Translate [192]. This automatic translation was then checked and adjusted by five different native speakers (where this included names of popular online stores).

3.3.2. RECRUITMENT AND PARTICIPANTS

Participants were recruited on the crowd-sourcing platform Prolific with the following screening criteria: Participants should reside in one of the five member states, be fluent in the respective official language, have an approval rate in prior studies of at least 95%, and completed at least five other submissions on the platform. We opened a separate survey invite for each age group to avoid bias towards younger cohorts, as Prolific samples tend to be younger [223]. In each age group, we also applied a quota for gender to get an even split. For the highest age brackets, we were not able to attain the full target sample size or an even gender split in all countries, as the number of available participants on Prolific for this target group was too small. We redistributed the remaining slots across the different age groups. The final sample consisted of 993 respondents and Table 3.2 depicts the demographic information.

Demographic		FR	DE	PO	ES	NL	Total
Age	18-24	39	41	48	39	43	210
	25-34	41	40	46	40	44	211
	35-44	43	41	43	42	45	214
	45-54	38	40	43	42	41	204
	55-64	25	29	14	33	20	121
	65+	11	9	2	3	8	33
Gender	Female	95	97	93	94	96	475
	Male	98	100	99	100	102	499
	Other response*	4	3	4	5	3	19
		197	200	196	199	201	993

Table 3.2: Demographic overview. *Other responses possible were "Non-binary/ Third gender" and "Prefer not to say"

We decided against screening for smart device ownership, given that the expectations of consumers without smart devices have legal relevance – they might be future users, or have their own particular reasons for not engaging with such products. Thus, we were able to compare expectations and perspectives between experienced and less experienced users of smart devices. We followed methodological literature for comparable experimental designs (i.e., factorial designs and choice experiments) to find sample sizes with sufficient statistical power for our planned device and country comparisons. [12] and [22] suggest a range of at least 5 to 50 observations per vignette as a general guideline for such designs with repeated measures (vignettes) per respondent and deploying multi-level regression for analysis. As we wanted to compare lifetime expectations of a total of 17 different device vignettes within each of the five countries, we aimed at 200 participants per country, as this would provide 35.3 responses on average per device per country, well above the minimum recommendation of five observations.

During data collection, we noticed that many respondents interpreted the survey questions about their own devices (Q2.1 - Q2.8) differently as intended by us, i.e., as asking about the general device category, not the particular device they were using right now (i.e., asking when they started using smartphones *in general*, instead of when they started their *currently used* smartphone). We thus added a clarification to Q2.2 to focus responses on the currently used devices, changed the wording of Q2.2 to Q2.8 accordingly, and then did a follow-up survey only with these specific questions with all respondents participating so far. We contacted them directly via Prolific and paid a higher hourly rate to incentivize participation. In total, 95% (552 of 579) did the survey with only the updated questions again. We conducted consistency checks against original responses and included an attention check. We discarded the responses from the remaining 5% when analyzing responses to Q2.1 to Q2.8, as their question interpretations might have been inconsistent with ours.

3.3.3. SURVEY PROCEDURE

The survey was advertised on Prolific as “*Your use of and perspectives on smart devices*” together with a short, high-level description of the task in the respective language. Participants could access the survey after agreeing to the informed consent. The survey was run on Qualtrics (hosted by the research institution) and structured into four sections, which we describe in the following.

Participants’ own smart device usage. To collect concrete, device-level data on respondents’ own smart device usage and duration, the survey began with questions about their own devices. The survey started with a general description of smart devices to provide context, especially for respondents without much experience with such devices. Respondents then indicated the types of smart devices they used and owned and which specific devices they currently used, followed up with more targeted questions about these specific devices, including the model/brand, the condition (new / previously owned), the purchase channel, when it was started being used and for how long looking ahead, and the software updates status (Q2.1 - Q2.8). We asked these follow-up questions only for a maximum of three of their devices (chosen randomly) to keep response time manageable. We also asked if respondents had recently stopped using any smart devices to find out more about their use time, device disposal, and the reasons for ending use (Q4 - Q4.4).

Eliciting lifetime expectations via device vignettes. Respondents were presented with vignettes about different devices to move from their own devices’ usage to eliciting their expectations regarding stereotypical device categories. We asked them to imagine they would buy such a device today and emphasized our concept of product lifetime (Q5). Then they were presented with one conventional and two smart device vignettes, both selected randomly from the respective list of devices (see Table 3.1). The first vignette was a non-smart device to set a baseline and avoid ordering effects by first showing smart devices, as this could easily lead to confusing the subsequent conventional device as being smart.

After each device vignette, participants answered the question: “*If you had to predict, for how many years do you expect such a device to last?*” (Q5.1.1) by entering a whole number in years. We opted for an open-ended response format to avoid potential biases introduced by scale anchors (*anchoring effect* [228]). For instance, an upper anchor such as “*More than 10 years*” might implicitly suggest that 10 years is a long lifespan, thereby lowering lifetime estimates [208]. To capture more contextual data for these numerical lifetime estimations, we also followed up with “*What aspects did you take into account when estimating the number of years?*” (Q5.1.2), where respondents could input their reasoning in an open text field.

Software update and security considerations. In this section, we were interested in respondents’ understanding of software support and security aspects for smart devices. As previous work implied that consumers did not fully understand the implications of unsupported smart devices [103], we asked how the end of update support would impact smart devices (Q7) and to what extent respondents had first-hand experience with this (Q7.1), followed by their expectations how long software support would last for the

two smart devices they saw earlier in the vignettes (Q8) to see how this estimate would compare to the device's lifetime expectation in Q5.1.1. We then explicitly introduced the concept of security for smart devices with a short explanation (Q9) to then measure respondents' security concerns (Q9.1), mitigation actions (Q9.2), and if a continued provision of security updates would lead them to use their smart devices for longer (Q9.3).

Demographics and survey conclusion After filling out some final demographic questions, respondents were thanked, debriefed, and sent back to Prolific with a completion code to receive payment. The survey included two attention checks (Q3 and Q6). The median completion time was 11.83 minutes. The full survey instrument is provided in Appendix B.1.

3.3.4. DATA ANALYSIS

We first checked the data for suspicious responses by reviewing attention checks, completion time, and spurious patterns like providing the same response continuously or nonsensical open text responses. All respondents who failed both attention checks were removed ($n = 5$). 64 respondents failed one of the two, and after manually reviewing their responses, we removed three due to failing at least one of our removal criteria; repeated scale responses, nonsensical text responses, inconsistencies in responses, such as having started using a device model that was not released yet at this time, or rapid completion times. We also checked the responses of the fastest (< 5 minutes) and slowest (> 1 hour) completion times, but did not find any indications of bogus answers (i.e., they passed both attention checks and gave legitimate text responses). Thus, response quality was generally high, as we only had to remove data from eight respondents.

Quantitative analysis To analyze the data, we first calculated descriptive statistics of survey responses to get an indication of patterns and trends as well as inferential statistics to detect statistically significant differences across responses or countries.

To assess which factors influenced respondents' expectations of devices' lifetimes in the vignettes, we ran a multi-level regression model with maximum-likelihood estimation, allowing us to test for significant differences between devices (both smart and conventional) and account for respondent-level characteristics, respondent-level variance, and model fit. Assumption checks for homoscedasticity, multicollinearity, overfitting, and influential datapoints indicated that lifetime expectations were not normally distributed and displayed a skew towards higher values (e.g., 20 years was not an uncommon response). We thus log-transformed the outcome variable *Lifetime*. For better interpretability, we transformed the resulting coefficient estimates to express percentage changes in the outcome variable (by exponentiation with $e^{\beta} - 1$).

After iteratively adding predictors to the model to assess if goodness of fit improved significantly, we concluded with the final model predicting *Lifetime* with *Device* (the different device vignettes), *Device Experience* (if the respondent used the smart device described in the vignette themselves), *Country*, an interaction between *Country* and *Smartness* (if the device was smart or conventional), and a random effect of respondent to account for individual differences between them. The number of smart device categories used, Age, and Gender showed no effect, as including those predictors did not

contribute significantly to the model's fit. The model is presented in Table B.1 in the Appendix.

Qualitative analysis Open-text answers were translated into English using the translation tool DeepL [62]. In our Ethical Considerations statement, we detail to how we protected respondents' sensitive data. To ensure quality, the native speakers who were also involved in the survey translation checked a random sample of the resulting translations. The automatic translation performed satisfactorily, with no native speaker reporting any translation errors. Translated text data was then analyzed using thematic analysis [32], where the primary researcher identified and assigned codes to the text answers. After coding a subset of the responses, the resulting initial codebook was discussed at regular intervals with co-authors and refined. After coding approximately one-third of the open text data, we reached theoretical saturation, where no new themes emerged [154], and the codebook was discussed one last time to finalize it. Previous codes were adjusted according to the final codebook, which is provided in Appendix B.4.

3.4. RESULTS

3.4.1. INDIVIDUAL USE OF SMART DEVICES

To address RQ (i) (*How do consumers use their smart devices, and for how long?*) and understand what is a reasonable use time for smart devices and how they are used and purchased, we collected data on survey participants' use of their own smart devices, which resulted in data on 2753 individual IoT devices currently used, and 398 devices used in the past. We also report on relevant country differences throughout this and the following subsections to address Research Question (iv) (*Are there differences among EU member states regarding consumers' smart device usage, expectations, and security and software support perceptions?*).

DEVICES CURRENTLY USED

On average, respondents indicated using five different smart device categories, with smartphones being the most common, followed by Smart TVs and Routers. Table 3.3 provides an overview of devices' popularity in the sample. Devices' most popular brands were all from international companies, who will need to comply to the CRA to sell these products in the EU. For some smart devices, especially routers, sensors, doorlocks, doorbells, and solar inverters, a sizable proportion of participants did not know the brand of their currently used device. We added smartphones as an option due to their widespread use and importance in the smart device ecosystem (e.g., due to companion apps).

To assess how long respondents used their devices, we shifted the level of analysis from device *categories* used to the *particular devices* participants were currently using. We asked if these devices were acquired in new or second-hand condition, as we could not control for the use duration of the previous owner when measuring device use times. This also provided us with valuable insights into the prevalence of the second-hand use of smart devices, as their lifecycle can span different users. The vast majority of smart devices (90%) were reported to have been in new condition, while only 9% were previously owned. The devices with the highest rate of second-hand use were smartphones (16%), smartwatches (11%), and smart TVs (9%). For the devices that were new when

Device Category	Number	Most common brands	% Unk.
Smartphone	984	Apple, Samsung, Xiaomi	0%
Smart TV	680	Samsung, LG, Sony	3%
Router	611	Fritzbox, TP-Link, Vodafone	37%
Connected printer	582	HP, Canon, Epson	7%
Smartwatch	503	Apple, Samsung, Xiaomi	5%
Media streaming	290	Google, Amazon, Xiaomi	14%
Smart lighting	281	Philips, Ikea, Xiaomi	27%
Smart speaker	281	Amazon, Google, Apple	4%
Robot vacuum cleaner	264	iRobot, Xiaomi, Roborock	11%
Smart security camera	192	Ring, Xiaomi, TP-Link/tapo	27%
Smart sensors	150	Eufy, Kidde, Philips	59%
Smart thermostat	111	Honeywell, Remeha, Tado	34%
Smart washing mach.	108	Samsung, LG, Beko	9%
Smart doorbell	88	Ring, Eufy	33%
Smart hub	48	Google, Philips	21%
WiFi solar inverter	48	Enphase, Solaredge	39%
Smart fridge	42	LG, Samsung, Bosch	5%
Smart doorlock	15	–	43%
Smart baby monitor	14	–	17%

Table 3.3: Prevalence and brands of smart device categories owned and used by respondents. If fewer than three brands are listed, there are too few observations for a pattern. % *Unk.* is the share of respondents not knowing their device's brand.

usage commenced, the majority were purchased in online stores (60%), while physical stores were less popular yet still a sizable amount (30%), highlighting the importance of considering both online and offline sales channels when disclosing support durations to consumers.

There were significant differences between countries regarding smart device purchase channels. Regarding second hand use ($X^2(4) = 28.39, p < 0.001$), the highest prevalence was observed in France with 14% of devices being acquired in used condition and the lowest in Spain (5%). Similarly, France had the highest rate of participants selling their previous devices onward (16%). For devices acquired new, the highest proportion of devices bought online was in Poland (69%) and The Netherlands (66%), while the countries with the highest rate of devices bought in physical stores were France (37%) and Spain (35%) ($X^2(8) = 33.47, p < 0.001$).

To gauge how long respondents use their current devices, we added the time since they started using them (Q2.2) and for how long they plan to keep using them (Q2.6). We consider this a measure of intended use duration, as it pertains to planned future use, which bears some level of uncertainty. Figure 3.1 depicts the distributions of past and intended use time across respondents' smart devices. Second-hand devices were not included. Thus, a first differentiation across smart devices in intended use times became evident. Smartwatches and smartphones were clearly skewed towards shorter use times

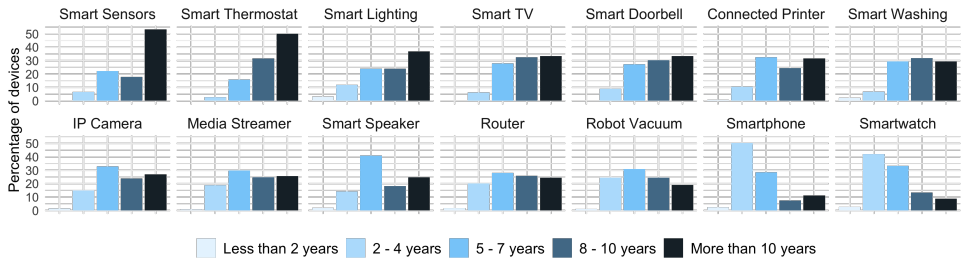


Figure 3.1: Sum of past use and intended future use duration across respondents' smart devices. Ordered by decreasing rate of durations of more than 10 years. Devices with less than 20 observations are not depicted.




Figure 3.2: Duration of device usage before stopped being used. Ordered by increasing rate of durations below 1 year. Devices with less than 20 observations are not depicted.

(i.e., mainly in the 2 - 4 years range), while Smart TVs, smart Lighting devices, or smart thermostats showed substantially longer use times on average. Indeed, we found that for the majority of devices except for smartphones and watches, a sizable number of users indicated more than a total of 10 years, indicating long intended use times in terms of smart product standards.

DEVICES RECENTLY STOPPED TO BE USED

As intended future use can be difficult for individuals to forecast, we also asked respondents about their recently stopped smart devices to assess their recent experience with a device they finished using (Q4.1 - Q4.4), independently whether it was replaced with a newer device or not. In total, respondents reported 398 smart devices they had recently stopped using. The most commonly stopped devices were smartphones ($n = 98$), smartwatches ($n = 81$), and robot vacuum cleaners ($n = 40$). Figure 3.2 presents the distributions of devices' use times (devices with less than 20 observations are not depicted) and shows that robot vacuum cleaners, smart lighting, smart speakers, and smartwatches were skewed towards shorter use times while smart TVs and printers were often used for longer. Smartphones were approximately normally distributed around a middle point of four years of use.

The most common reasons why respondents quit using their previous devices were wanting a newer device ($n = 97$, 20% of all reasons provided) and experiencing flaws with the old device ($n = 93$, 19%). Reasons implying a more terminal cause for stopping device use were less common, with 57 (12%) instances of devices stopping working altogether, and 34 (8%) in which incompatibility with other devices or services was the

reason. The end of update support ($n = 23$, 5%) and privacy ($n = 26$, 5%) and security concerns ($n = 16$, 3%) were the least common reasons to stop device use.

Similar motivations for device replacement also emerged for respondents who indicated that they already planned to only continue using their current devices for less than a year into the future, in Q2.6. When asked about their motivations, the most prominent reasons were experiencing issues and flaws ($n = 62$, 34%) and desire for a newer device ($n = 48$, 27%). Only 14 participants indicated the end of update support as a reason, and even fewer were due to security and privacy concerns. Thus, users' preference for better alternatives and experiencing functional issues with the current product were the main drivers to stop using their smart devices, rather than update support or security concerns.

Considering what happened to the 398 discontinued devices, the vast majority of respondents indicated they still have them while not actually using them ($n = 190$, 48%). Gifting it to others ($n = 54$, 14%), selling it ($n = 48$, 12%), or discarding it ($n = 33$, 8%) were substantially less likely responses, implying that handing devices on (i.e., gifting or selling it) and subsequent re-use were not common while the vast majority of devices appear to remain with their owner's without being used.

3.4.2. LIFETIME EXPECTATIONS OF SMART DEVICES

DEVICE LEVEL DIFFERENCES

As a respondent's use time of their own device does not necessarily correspond to its expected lifetime (e.g., they could opt to replace and sell it after a short while), we pivot to what respondents expected for how long different device types would generally last for whoever was using them by presenting vignettes of these products. To answer RQ (ii) (*How long do consumers expect different smart device categories to last, and which factors influence these expectations?*) we aimed at learning for how many years participants predicted a range of different smart and conventional device types would last, as elaborated in Section 3.3.1. We supplement these results with qualitative commentary given by respondents when asked what they took into account when estimating devices' lifespans. Table B.2 in the Appendix provides the five most commonly mentioned factors per device type taken into account.

Figure 3.3 presents the expected lifetimes provided for different smart and conventional device categories. The highest expected lifetimes had solar inverters (Conventional: $M = 12.65$ years, $SD = 6.60$, Smart: $M = 10.75$ years, $SD = 6.81$) and thermostats (Conventional: $M = 11.37$ years, $SD = 5.68$, Smart: $M = 10.35$ years, $SD = 6.29$). For solar inverters, respondents commonly considered their use case and high expected price as factors (e.g., "*It is an expensive product and should last according to its cost.*" PID678), and for thermostats their usage or perceived lower complexity (e.g., "*That it is a simple device*" PID685), as seen in Table B.2.

Devices with the lowest expected lifetimes were consistently smart. The lowest lifespans were expected for smartwatches ($M = 4.84$ years, $SD = 1.96$) and robot vacuum cleaners ($M = 5.74$ years, $SD = 2.51$). For both, their continuous use was commonly considered (e.g., a robot vacuum cleaning the floor regularly, a smartwatch worn continuously), while for robot vacuums, environmental factors like animal hair often came up. For smartwatches, however, factors more related to smart device market dynamics




Figure 3.3: Expected lifetimes for different smart and conventional device categories ordered by mean expected lifetime. Responses were given as whole numbers. The dashed horizontal line denotes the CRA's five-year baseline, and asterisks significant differences to the reference level *Smart speaker*.

were mostly considered: fast innovation cycles, software updates, degrading batteries, and planned obsolescence (e.g., “*Technology changes so fast, it is likely to be obsolete in 5 years*”, PID80). Expectations also generally showed less variability for these devices, as evidenced by lower standard deviations and error bars in Figure 3.3, indicating that expectations reached a stronger consensus among respondents for those products. This was likely due to respondents’ more common experience with smartwatches and robot vacuums than for solar inverters and thermostats, as evidenced in their respective prevalence of use (Table 3.3).

The multi-level regression analysis (Table B.1 in the Appendix) revealed that differences in expected lifetimes across devices were statistically significant. Taking smart speakers as the reference level (due to their central role in many smart home set-ups and their “important” categorization in the CRA), we found the largest differences to conventional solar inverters (expected to last 56.90% longer) and thermostats (52.70% longer), but also significant differences to other smart devices such as smart washing machines (expected to last 23.20% longer) or smartwatches (expected to last 22.10% shorter). Significant differences thus correspond to the visual trend in Figure 3.3, where also significant differences are highlighted.

We also found that smart devices had a consistently shorter expected lifetime than their conventional counterparts when considered as a factor. Most smart devices (except for smoke detectors and printers, with a marginal difference) had a lower average reported lifetime, an effect that was statistically significant (-17.1% , $t(2972) = -9.82$, $p < 0.001$). We identified several drivers for this, such as the added complexity of smart features (“*Adding [smart] functions risks making them even more susceptible to problems.*” PID564), software updates, or potential incompatibility, but also some more complex dependencies were taken into account, like “*The company selling this machine will*

go bankrupt, and we won't be able to install the application on updated OS versions." (PID21).

As can also be derived from Table B.2, devices' *Usage* and *Device type* were generally among the most common factors considered by participants when estimating lifespans across all devices. While static device qualities such as the type of device ("Washing machines last a long time") can be seen as more straightforward heuristic for the lifespan, usage ("How often the washing machine runs") is less deterministic, as this depends on wear and tear of components, how frequently and intensively the user decides to use the device, and if it is used as intended or not.

INDIVIDUAL DIFFERENCES

Lifetime expectations were also strongly affected by respondent-level factors besides device-related aspects. The regression model fit coefficients provided in Table B.1 suggest that when accounting for respondent-level variation, the model fit (Conditional R^2 : 46.40%) increased substantially in contrast to only considering the device, country, or own experience with the device (Marginal R^2 : 16.60%). This highlights the substantial impact of individual differences and the subjective nature of lifespan expectations for products. Some respondents consistently expected longer lifespans across devices, while others had a tendency to expect shorter lifespans, indicating systematic variation in personal perceptions.

Also personal experience with the smart device affected expectations - respondents who used the same device type themselves as described in the vignette at hand expected longer lifespans on average than participants who did not (8.0% longer, $p < 0.01$), possibly due to familiarity with the device or increased optimism due to personal investment in it. However, we found no effect of experience with smart devices in general (i.e., the number of smart device types the participant was using), implying that the user's experience with the same device category was a stronger predictor for expected lifetimes.

We also found significant differences between countries in how long respondents expected devices to last. Taking Spain as the reference, respondents in Poland generally expected devices to last significantly shorter (14.1% less, $p < 0.001$), while participants in The Netherlands and Germany had more optimistic expectations (9.2% and 8.6% longer, respectively). However, an interaction between country and the smartness of the product indicated that in these two countries, respondents expected smart devices to last significantly shorter than conventional devices compared to Poland (Germany: 8.5% shorter, The Netherlands: 12.2% shorter). Thus, participants in Germany and The Netherlands were comparatively more skeptical about the lifespan of smart devices. Age and Gender had no effect on expected lifetime.

3.4.3. PERCEPTIONS OF SOFTWARE UPDATES AND SECURITY OVER DEVICES' LIFETIME

As the previous analyses focus on smart devices' use times and expected lifetimes to measure what constitutes a reasonable lifetime, we now focus specifically on the role of software updates and security aspects over devices' lifespan to answer RQ: *(iii) How do consumers perceive security and software update support over smart devices' lifespans?*

UPDATE EXPECTATIONS AND EXPERIENCES

In contrast to the static nature of most conventional, non-smart devices, smart devices can be under continuous modification due to software updates. We were thus interested in understanding how consumers perceived update-induced changes to devices over their lifespan and how their subsequent end would affect the user's experience with the product.

More than half of respondents ($n = 533$, 54%,) expected the software update support to correspond with the device's full lifespan, as they provided the same value for both (Q5.1.1 and Q8). When directly asked for their expectations on how the end of update support would impact smart devices' use, no new features ($n = 853$, 86% of participants) and incompatibility with other devices and services ($n = 752$, 76%) were the most commonly expected changes, while almost no participant expected no changes (Figure 3.4). Interestingly, both security-related answer options (no more fixes to vulnerabilities and no more risk monitoring for the device) were also commonly expected ($n = 714$ and $n = 705$, respectively), despite the survey not mentioning security or privacy up to this point and answer order randomization. Thus, approximately 70% of the respondents were indeed aware of the security implications that the end of update support brings with it.

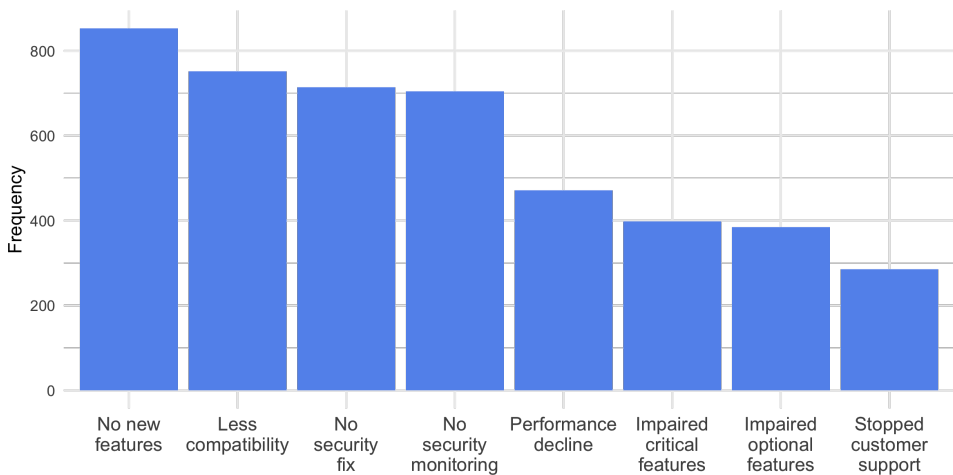


Figure 3.4: Changes to smart devices after software support ends as expected by respondents (in %). The order of the options was randomized. Multiple choices were possible. The responses “*I don't expect any changes*”, “*Not sure*”, and “*Other*” were below 1% and omitted from this plot.

To check respondents' own experiences with end-of-support-induced changes, we also asked if they had experienced their previously provided expected changes with their own devices. Figure 3.5 shows that no new features, a declining device performance, and less compatibility with other devices and services had the highest proportion of respondents experiencing this. We also saw the potential role of ended support for product obsolescence: Among the participants who expected critical features to stop working at this point, around half actually experienced this, so their devices' functionality was




Figure 3.5: Respondents' own experiences with changes to their devices after the end of update support. They were asked: "Did you experience any of these changes with your devices yourself?"

essentially made obsolete after not receiving updates.

While such functional impacts on the device were thus quite apparent, security-related changes were more subtle and imperceptible for users. Both security-related answer options (no more fixes to vulnerabilities and No more risk monitoring for the device) had the highest rates of respondents indicating not experiencing them or not being sure, while they both were frequently expected as a common change after the end of support (see Figure 3.4), highlighting how opaque security (and its absence) can be for smart device users.

Related to this imperceptibility of updates, Figure 3.6 shows substantial differences in the current update support status across the smart devices used by respondents. For smartphones, the vast majority indicated the device to still be receiving updates, while for many other smart devices a substantial proportion of participants reported not knowing if it still was. For instance, for smart sensors and lighting products, at least half of respondents indicated they did not know if the device still received updates. Indeed, this uncertainty of the device's update status was much more prevalent than unsupported devices still being used (with robot vacuum cleaners having the highest proportion of unsupported devices). Thus, while the majority of respondents' smart devices were reported to still be supported, it also illustrates the challenge for users to actually *know* whether their smart devices still receive updates, especially for devices with limited user interfaces like sensors, lighting, doorbells, or routers.

SECURITY AND PRIVACY PERCEPTIONS AND BEHAVIORS

As a majority of respondents were aware of the security implications of the end of support, although often imperceptible, we now focus on their security and privacy-related perceptions towards smart devices.

Figure 3.7 suggests that when prompted for security-related concerns about their own smart devices, respondents mostly indicated to be concerned about their personal data ("My personal data being accessed by unauthorized individuals" = 57% of respon-




Figure 3.6: Responses to the question: “Does the particular device you are currently using still receive software updates, as far as you know?” Devices with less than 20 observations were not included.

dents, “More personal data collected from the device than I expect” = 56%, “My personal data being shared with third parties by the manufacturer without my consent.” = 52%). Two of these were actually privacy-related concerns, with the device’s vendor as the “perpetrator”. The device being directly attacked and leveraged maliciously was also a common concern (50%), while the device being damaged by an attack (30%) or insufficient customer support in case of an attack (23%) were expected less often. Not being concerned was the least selected response (15%).

To address such concerns, respondents indicated to apply varying mitigation actions, as depicted in Figure 3.8. The importance of updates became evident, as installing the latest software update was the most prominent mitigation behavior and reported by 56% of respondents ($n = 552$). Stopping to use the device altogether after software support ends was the least common behavior ($n = 121$), yet still selected by approximately 12% respondents. This was in line with our previous finding that the end of software support or concerns about privacy or security were not prominent reasons to stop using a device.

When asked if respondents would use their smart devices for longer if they continued to receive *security* updates (Q9.3), the vast majority strongly agreed or agreed (78%). This seemed to be a contradiction at first, as the end of update support was not a common reason for respondents to stop using their devices nor a trigger to stop using the device to address security concerns. We interpret this finding such that as longer patching support will give reassurance to users that they can safely *continue* using their devices if they want to, its end was not a strong enough trigger for users to *quit* using them. We also note

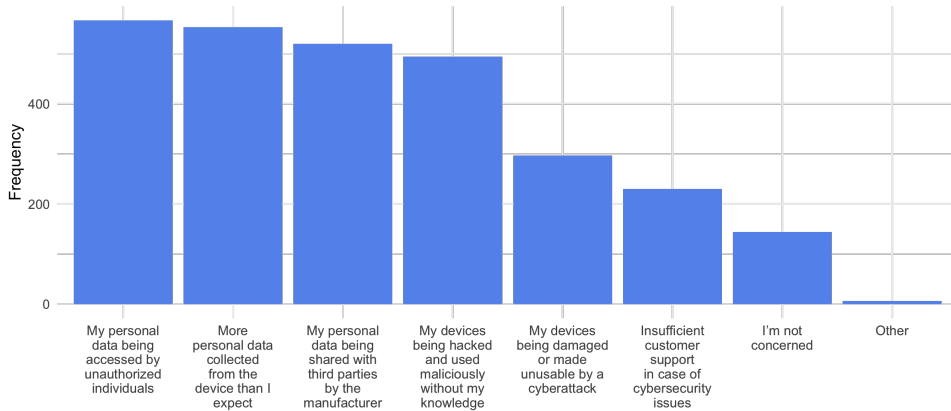


Figure 3.7: Respondents' concerns relating to security and privacy risks with their smart devices. The order of the options was randomized. Multiple choices were possible.

the importance of the device's brand and its associated trustworthiness for respondents to avoid security risks (53%, $n = 528$).

We also found significant country differences in how participants reported to address their concerns ($X^2(24) = 38.03$, $p < 0.05$). Installing the latest updates was reported the most common mitigation in the Netherlands (24%) and Germany (23%), but less so in Poland (19%), where, as well as in Spain, being cautious about how using the device and choosing trustworthy brands were more popular.

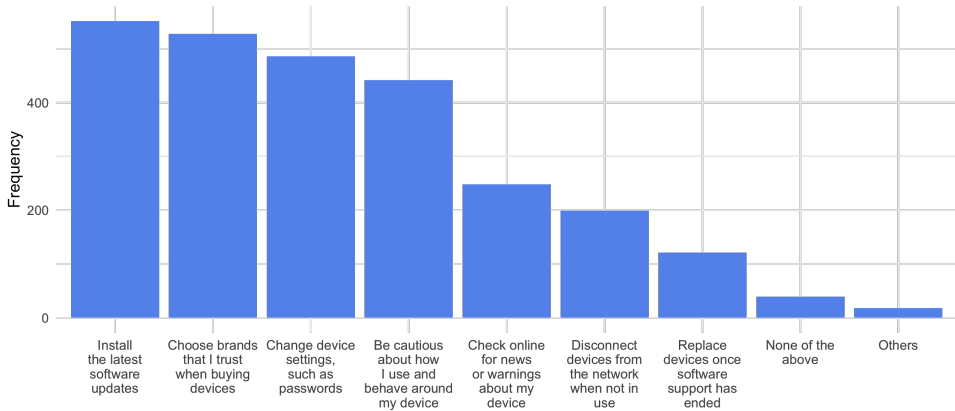


Figure 3.8: Respondents' mitigation behaviors relating to their security and privacy concerns with their smart devices. The order of the options was randomized. Multiple choices were possible.

3.5. DISCUSSION

In this survey study ($n = 993$), we empirically explored consumers' usage and expectations of smart devices and their lifespans in the EU, to address the uncertainty around the abstract legal concept of reasonable user expectations. The Cyber Resilience Act [243] aims at converging security support duration with smart device lifetimes, and requires manufacturers to consider reasonable user expectations when defining this period – a policy with global ramifications for the smart device market. For instance, practically all brands of respondents' smart devices were from large, globally active manufacturers, who will have to comply and could thus also uphold support in other markets.

3.5.1. DEVICE LIFETIMES

We found that all smart devices included in this survey (except smartwatches) were on average expected to last longer than the CRA's minimum support period of five years, despite many participants considering factors artificially reducing product lifetime, such as a premature replacement for newer alternatives or planned obsolescence. Smart products were also commonly used by respondents for a period longer than this, with use times well above 10 years common for some devices. Surprisingly, a majority (54%) of participants actually expected the update support period to correspond to the smart device's full lifetime, which contrasts with current practices in the smart device market [81, 189, 196], but highlights the importance of aligning support periods to be closer to device lifetime to meet expectations.

Previous work found that device re-use is a common practice, with the first owners often gifting or selling their devices onwards [109, 140]. We found this was present but not substantial in our sample. 9% of respondents' own devices were second-hand, and selling or gifting a previously used device was by far not as common as retiring it at home without using it. Still, second-hand use or repair [40, 248] are conceivable practices, perhaps due to economic reasons (e.g., for low-income users [128]) or sustainability consid-

erations, increasing device use times. It is unclear to what degree second-hand use and repairability will be factored in when manufacturers, market surveillance authorities, or courts determine expected use time and support period.

3.5.2. SOFTWARE UPDATES AND SECURITY OVER DEVICES' LIFESPAN

“Smart” factors such as software updates, incompatibility, or added complexity led to smart devices being expected to have shorter lifetimes than conventional products (Section 3.4.2). However, the central role of software updates for security was often seen. In contrast to a commonly reported unawareness or uncertainty of updates among users in previous research [45, 75, 103, 104], we found that respondents generally exhibited awareness of updates' security implications. 70% expected no more security reassurances after update support ended, and installing the latest updates was reported as the most common mitigation to address security concerns. As there are some differences in our study's demographics in contrast to this previous work, with participants in [103, 104] being US-based smart home users with limited security experience, [75] studying young users of widely used software, and [45] practitioners at UK router vendors indirectly describing their users, it raises the question if this effect was due to differences in samples and or if this awareness might actually be increasing over time.

However, while an association between end of update support and security implications was arguably observable in our study, security concerns were weak triggers for respondents to stop using a device. More attractive alternatives or declining functional performance were the prevalent drivers. This was likely amplified by the imperceptibility of updates and security measures. For many of their own devices, respondents did not know if the device was still supported, or did not experience or know about security support stopping.

Thus, a key challenge with prolonged support duration is ensuring visibility to users. Currently, indicators of software versioning and update status are often unclear or inconsistent across smart devices (as observed in our work and in, e.g., [102–104]), which often lack a user interface. It is challenging for users to verify if a manufacturer is maintaining security updates as promised, potentially leading to a form of learned helplessness akin to delegation of security responsibilities [68], where users must rely on manufacturers' assurances without a straightforward means of verification for themselves. While the CRA sets requirements on update delivery and notifications (to be automatic by default, with notification channels to users available, e.g., to opt-out), it remains an open question through which mechanisms - if any - users can actively engage and hold manufacturers accountable for ongoing security support, comparable to other regulatory frameworks such as GDPR [231], where subject access requests can be submitted to obtain data about themselves.

As the CRA is still in a nascent stage, guidance about support periods will be crucial in the coming years before the market converges on expected support times for “products with digital elements” in all their forms. Guidance will reduce uncertainty and avoid courts having to eventually make the case for what constitutes a reasonable use time and which elements of a product determine this. Our work contributes to this ongoing conversation by empirically assessing device use times and expected lifetimes for a range of different smart product categories and what factors consumers take into account, as

information valuable for manufacturers, market authorities, and the EU administrative cooperation group (ADCO), who will publish guidance on support times in the coming years.

3.5.3. RECOMMENDATIONS

Here we discuss initial recommendations emerging from our work.

3

Manufacturers: For manufacturers, users' expectations take center stage in product conformity with the CRA; manufacturers will now need to document how they consider these expectations when determining support periods.

Going beyond device-internal usage data. Our work demonstrates that for a meaningful understanding of users' expectations, getting in contact with them is essential, as relying exclusively on internal device usage data (i.e., device logs how long it is used) might be insufficient to account for expectations, as other contributing factors emerged from our data. For instance, respondents differentiated between smart and non-smart products and considered factors such as the device's price, perceived quality, or environmental factors of usage. Also, the duration how long device data can be collected is arguably determined by the manufacturer, not by expectations. We thus recommend that manufacturers leverage sales channels as well as market and user research capabilities to learn about and document their customers' lifetime expectations for the respective product categories. As our data exhibited substantial individual and country differences, we advise applying a sampling strategy that considers countries, user demographics, and sufficient sample sizes.

Reducing user uncertainty. Furthermore, our results suggest that ongoing security support for smart devices is often indistinct to users, as evidenced by high proportions of users not being sure about the update status of their own devices (Figure 3.6) and experiencing uncertainty with respect to the end of update support (Figure 3.5). Thus, tangible factors demonstrating competence to provide support should be made salient by manufacturers at the time of device adoption and over the device's lifetime, possibly woven into marketing, product disclosure, and easy-to-access support status indications of the product (e.g., via a companion app). As we also found that choosing trusted brands was a major way for respondents to navigate security concerns (Figure 3.8) and a clear preference for longer support periods (as also suggested by previous work [73, 160, 164]), a stance of active security support has marketing and brand-building potential.

Policymakers and market surveillance: Approaching consumer expectations. Market surveillance authorities in EU member states should note that expected lifetimes and actual use times of consumer smart devices were commonly found to be higher than the CRA's baseline of five years (Figure 3.1, 3.2 and 3.3). The EU administrative authorities (ADCO) can publish guidance on reasonable support durations for different product categories, or, if "data suggests inadequate support periods for specific categories of products" (Article 52(16) & Recital 62) issue recommendations to market surveillance authorities to focus on those products or to the EU Commission to delegate acts and define minimum support durations. Our data should thus be considered as a first indication that several smart product categories considered in this study require delineation from the baseline of five years as suggested in the CRA. Additionally, they should be aware

of a potentially self-defeating dynamic for consumers regarding expected smart device lifetimes: Smart versions of a product category were expected to last shorter on average than conventional counterparts, and participants considered factors they are accustomed to from the current smart device market; quick replacement, emerging incompatibility with other products, or planned obsolescence. As these are factors artificially reducing use times and thus not justifying a product's definite end of life (and smart devices being repaired, resold, and reused, as we also found in our study), policymakers and market authorities should be cautious to consider such dynamics when user expectations are measured or defined.

Providing assurances to users. The CRA's support duration provisions will likely lead to longer support with security updates to smart devices, exposing them to continuous modification. We found that participants experienced phenomena with their devices over their use time that may or may not relate to (the end of) security updates, like declining performance, less compatibility, impaired features, or simply being unsure (Figure 3.5). Thus, policymakers and market authorities should consider defining assurances for users over the devices' lifetime that manufacturers should provide. This could include explicitly calling out potential "side effects" how the update could affect the device's behavior via notifications accompanying updates. As the CRA also requires security updates to be published separately from functionality updates where technically feasible (Annex I, Part II (2)), we recommend market authorities to explicitly consider this legal instrument to reduce such in-transparency for users.

3.5.4. LIMITATIONS

A survey relies on self-reporting, and for some responses, such as whether longer security support leads to prolonged device use, it remains uncertain if this translates to actual behavior. Additionally, the survey topic and some questions were complex. For example, recalling the year a device was first used or whether it still receives updates could be difficult for respondents. We did not filter for smart device ownership, and while most respondents used at least one device, concepts like smartness, updates, and security could be challenging. To mitigate this, we used simple language, refined the survey through iterations, conducted a pilot study, and provided definitions and "I don't know" options.

Sampling and generalizability are further limitations. Our sample is not representative of the general EU population or the included countries, and we only studied a subset of product categories falling under the CRA. As a common limitation in such online surveys, participants on Prolific tend to be younger and more technologically knowledgeable than the general population [223], leading to an underrepresentation in our study of those over the age of 55, despite aiming for a balance across age groups and genders. However, even with a sample of participants with an arguably higher technological literacy, we observed high levels of uncertainty towards participants' awareness of their devices' support (e.g., Figure 3.6). Resource constraints also limited our survey delivery to a subset of five member states (with Section 3.3 explaining our rationale). We thus invite future work to extend our methodology to other countries and products with digital elements, as the CRA will apply to a plethora of products, not just consumer smart devices.

Lastly, many factors can influence a device's expected lifespan, including usage, en-

vironment, quality, and price. Our vignettes focused only on device type, exploring additional factors through qualitative responses. Future work could adopt more complex vignette designs incorporating variables such as brand, price, and usage.

3.6. CONCLUSION

The upcoming Cyber Resilience Act will increase the responsibilities of manufacturers of smart consumer devices and market surveillance authorities in the EU to determine and uphold the update support periods of smart products. In contrast to other legislation, where disclosure on support periods via product labels is prominent [80, 229], the CRA will require manufacturers to actually uphold a certain support period. In contrast to the previous status quo, this will involve considering users, as reasonable user expectations regarding the product's lifetime will have to be considered when determining the support period.

In this work, we empirically assess user expectations and behaviors regarding smart devices' lifetimes, support durations, and security considerations to provide insights into this abstract legal notion. We find that expectations vary for different (smart) device categories, with smart devices generally being expected to last shorter than their conventional counterparts. We also find that a prolonged support duration harmonizes with users' expectations.

4

SECURING CONNECTED MEDICAL DEVICES: PATCHING PRACTICES AND EXPECTATIONS AT ORGANIZATIONS

Medical devices become increasingly connected and thus require security measures to ensure patient safety and data protection. However, such connected medical devices are often reported to lack basic security and to run on unpatched and outdated software. Thus, there is an increasing push to deliver security patches faster and more regularly to devices in the field. In this work, we empirically study current practices of patching connected medical devices by conducting 23 semi-structured interviews with participants from nine healthcare delivery organizations (HDOs) and three medical device manufacturers, also capturing data on actual updating practices for 25 specific medical devices. We find that delivering software updates to medical devices is an laborious and costly process for HDOs and manufacturers, as operational demands for medical use and an increasing need for infrastructure management put significant strain on involved stakeholders, thus rendering it questionable if conventional security patching will actually work in the healthcare sector without overwhelming it operationally and financially.

This chapter has been published as: **Kustosch, L.F.**, Gañán, C.H., van Eeten, M.J.G., & Parkin, S.E. (2025). "Patching Up: Stakeholder Experiences of Security Updates for Connected Medical Devices". In *Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)*.

4.1. INTRODUCTION

Healthcare Delivery Organizations (HDOs), such as hospitals, clinics, and practices, operate with an increasing number of medical devices connected to their network. Think of imaging, patient monitoring, or surgery equipment. A key measure to secure any networked device from attacks is to regularly provide software updates to address vulnerabilities. However, it has been repeatedly reported that a plethora of connected medical devices remain outdated and vulnerable. The prevalence and risks of unpatched and outdated medical devices have been raised by governmental actors such as the FBI [179], industry reports [37, 56, 115], and academic studies [100].

The security risks of connected medical equipment have led to a regulatory shift, such as the PATCH Act, which mandates the U.S. Food and Drug Administration (FDA) to require manufacturers to continuously release security patches over the lifetime of medical devices [113]. Similarly, the E.U. Medical Device Regulation (MDR) requires manufacturers to respond to security risks with potential patient safety implications (e.g., with a security patch) and upcoming regulations in Japan require manufacturers to follow the IEC 81001-5-1 standard, which also defines a continuous risk management and patch release process for medical equipment. The regulatory pressure on manufacturers to release timely patches only makes sense if the HDOs deploy them, which would lead to an increased patching frequency of medical devices in the upcoming years.

It remains unclear, however, how faster patching cycles for medical devices interact with the operational status quo in hospitals. Prior studies [67, 100, 120, 127, 217] have interviewed IT and health professionals about the broader challenges of providing cybersecurity in HDOs. In most cases, patching is only mentioned in passing, if at all. That said, this line of research does consistently observe a difficult tension between security requirements and the operational pressures to deliver health services. Dissanayake et al.[67] did study patching in healthcare, but for conventional IT, not medical devices. The closest related work to our study is [100]. The authors interviewed eight IT professionals, also covering challenges of patching medical devices. However, the IT department is usually not actually deploying the patches to medical devices. That is typically done in the medical departments by so-called biomedical engineers. So the study provides an indirect view of IT professionals on the updating of medical devices.

We present the first study on the security updates of medical devices based on interviews ($n = 20$) with biomedical engineers and other professionals who actually control the devices and their patches. These interviews were conducted in 9 HDOs in the Netherlands, Italy and the U.K., typically responsible for thousands of connected medical devices. We complement these observations with interviews with three leading device manufacturers ($n = 3$) since patching actions might be conducted by the manufacturer's field technicians.

The purpose of this study is to understand how the process of updating connected medical devices in their operational environment at HDOs is implemented and managed, and what this process translates to with regard to security, patient safety, and operational effort for the involved stakeholders. We pursue two research questions: (i) How are connected medical devices patched within their operational environment at HDOs? and (ii) What kind of challenges do HDOs and medical device manufacturers encounter during this process and how are they mitigated?

In short, we make the following contributions:

- We provide novel evidence on the patching of medical devices in hospitals and find that around half of the HDOs try to patch devices as much as possible by themselves, while the others mostly outsource it to manufacturers or third-party providers. Rarely were patches deployed remotely, even though this option often exists. No HDO was able to install patches for all devices by themselves.
- No HDOs tracked vulnerabilities and all available software updates for their devices. They wait to be alerted by manufacturers or authorities. The frequency of update deployment varied greatly across device types and organizations, from once every 3-4 years to every 2 months. A rough mean for the examples we discussed was around once per year. In many cases, it was unclear to our respondents if their patching kept up with the release cycle of patches.
- We extend prior work on the security of connected medical devices by highlighting various factors that impact patching, such as cost. In some cases, patches are bundled with software updates that need to be paid for. Two participants mentioned that applying a single update on a single device would cost around 10,000 euros. Even when patches are free, the services of manufacturer field technicians are “expensive”.

In the next section, we introduce the background and related work for our study, including the regulatory environment. After that, we turn to the methodology, our results, discussion, and conclusions.

4.2. BACKGROUND AND RELATED WORK

This work focuses on physical connected medical devices specifically used for patient care at HDOs. We define them as cyber-physical systems used for hospital patient care, such as monitoring, diagnosis, surgery, and/ or drug delivery that are equipped with network capabilities. Thus, this study does not focus on medical equipment for patients' homes, more consumer-grade health-related IoT devices (such as wearables), or software-as-a-medical-device (SaaSMD).

4.2.1. IT SECURITY IN THE MEDICAL DOMAIN

Previous work studying IT security at HDOs frequently mentions the increasing connectivity of medical devices and end-point complexity as a significant risk factor increasing healthcare providers' vulnerability [7, 52, 120]. Security updates are also regarded as important for the cyber-resilience of HDOs [52]. A range of proof-of-concept attacks on connected medical devices have been demonstrated, such as hacking an insulin pump [186] or accessing a hospital's picture archiving and communication system via a connected CT scanner [159]. Medical staff may also have difficulty in detecting manipulated readings of a connected patient-monitoring system [258].

However, empirical work on how such medical IoT devices are actually deployed in their operational environment, how they are secured, and the organizational practices around them, is scarce. Previous work, via HDO scans or Shodan, has identified vulnerable and misconfigured connected medical devices [69, 151, 247]. Prior work has not

considered if and how devices are patched, and how affected HDOs manage related risks organizationally.

One of the few empirical accounts [100] reported severe challenges, such as significant delays in patching, a disorganized process that varies across devices, vendors, and hospitals, and staff's uncertainty concerning patch prioritization and timing. Coventry et al. [53] reported on hospital staff not feeling adequately prepared for the security implications of connected medical devices and the general diffusion of responsibility outside of typical biomedical or IT departments.

4.2.2. PATCHING PRACTICES

Although there is prior research of patching behaviours for end-users [77, 104, 150, 246], and systems managers in organizations [26, 64, 121, 138, 225], empirical literature on patching IoT devices within organizations is limited.

Li et al. [138] studied how system administrators implement and manage the patching process within organizations. Challenges were noted, such as difficulties in determining the availability of patches, bug fixes and security updates competing for prioritization, incremental testing and roll-out of updates, and updates introducing problems of their own. Other studies about patching practices from the system administrator perspective reach similar conclusions. [26, 67, 225].

Building on this work, Dissanayake et al. studied sociotechnical factors of the patching process within the healthcare sector, in a series of studies with a governmental health services agency and an IT service provider in Australia [65–67]. The participating organizations provided (security) updates to customers' IT server infrastructure, which could also be running hospital applications, such as Electronic Medical Records (EMR). The authors report on struggles with coordinating the patching process across varying departments, customers such as hospitals (e.g., due to the resulting system downtime), and medical software vendors, often resulting in delays. Further challenges include technical dependencies and compatibility issues with existing hard- and software (including outdated OS systems), and the mental overload for system administrators, who have to manage an increasing number of configuration options, patch releases, and software versions in a heterogeneous IT environment.

While previous work provides important insights into the updating practices of conventional IT (such as servers and workstation PCs) from the perspective of IT experts such as system administrators, it does not study connected medical devices nor HDOs' perspective on how this critical infrastructure is managed, secured, and updated.

4.2.3. REGULATORY LANDSCAPE OF MEDICAL IOT DEVICES

Medical devices are heavily regulated due to patient safety, which also impacts security implementations and the software updating process. Here we consider a selection of key regulations for connected medical devices from major markets. Medical device manufacturers often harmonize their processes to comply with most regulations worldwide. Thus, the Food and Drug Administration's (FDA) rules from the USA usually also apply to medical devices being developed and sold in other markets.

Depending on the medical device class and associated risk level, manufacturers have to provide documentation to the FDA and get premarket approval for devices considered

high-risk. The evolving security risk landscape for connected medical devices has led to the introduction of the PATCH Act, which defines new cybersecurity requirements for connected medical devices to be enforced by the FDA[2]. From October 2023 onward, to receive premarket approval, medical device manufacturers have to demonstrate to meet these requirements, which include postmarket surveillance of security vulnerabilities and having processes in place to release security patches on a '*reasonably justified regular cycle*', and for critical vulnerabilities, '*as soon as possible*'. This does not apply to devices already on the market unless any change to the device would require another premarket submission. Previously, the FDA provided non-binding security-related guidelines for manufacturers [112, 114].

In the EU, another influential market for devices, connected medical devices have to comply to the Medical Device Regulation (MDR) [232]. For a medical device to receive the CE label and be sold within the EU, it is assessed for adherence to the MDR's requirements by a notified body. Furthermore, the NIS2 Directive [240] directs member states to ensure that operators of critical infrastructure like HDOs take appropriate security measures, such as adopting cyber hygiene practices like software updates, while the GDPR [231] establishes requirements on data protection.

According to the MDR, medical device manufacturers need to ensure patient safety over the device's lifespan. Thus, security risks that could impact patient safety have to be resolved and changes to the device's hard- or software have to be validated to ensure continued safety. Thus, for any software update along the device's technology stack (e.g., an OS security update), the manufacturer needs to go through a validation process, which impacts the update release time. The operator can only install a software update on a medical device that has been validated by the manufacturer.

In the UK, the Medicines and Healthcare Products Regulatory Agency (MHRA) regulates the market and medical devices need to comply to the UK Medical Devices Regulation 2002[96]. The National Health Service (NHS) provides cybersecurity guidelines for HDOs (e.g., [167]). Regulations in the UK are currently in transition, with EU CE-marked devices still being accepted in the coming years, yet future regulations are in development[156].

A recent regulatory push comes from Japan, where medical device manufacturers are required from April 2024 onward to continuously improve devices' software security with patches according to IEC 81001-5-1 [119], regardless of an acute critical risk[98]. Thus, regulators increasingly recognize the importance of connected medical device security, and a general trend towards more rules for frequent and timely patch releases is observed.

4.3. METHODOLOGY

To explore patching practices of connected medical devices empirically, we conducted semi-structured interviews with 20 stakeholders at HDOs and three product security experts from three different major medical device manufacturers between July 2023 and January 2024. All participants were involved in the patching process of medical devices in some capacity, which allowed us to understand the process more holistically.

During the interviews, we also probed for details on the patching process of actual devices and thereby collected 25 cases of varying updating processes. This allowed us

to collect rich data on the actual practice of patching and its inherent variability across medical devices, operational contexts, and organizational structures, painting a picture of a heterogeneous, complex, and, at times, ad-hoc process.

4.3.1. RECRUITMENT AND PARTICIPANTS

As we expected general IT and updating practices to differ across HDOs and devices, we collected data from various HDOs to get a broader sample. The majority of the HDOs were recruited in The Netherlands as the researchers' country of residence, yet we also included HDOs from the UK and Italy to capture country variability. Medical device manufacturers were also European, two of which were headquartered in Germany and one in The Netherlands.

We recruited participants who are involved in the patching process of medical devices in some capacity, such that they either (i) are included in the decision-making around software updating processes, (ii) implement and/ or roll out updates, or (iii) are involved with product security at manufacturers.

To recruit a sample from this hard-to-reach population [78], we first leveraged our professional network and research project consortium to reach out to stakeholders working in the healthcare sector as initial points of contact. With these professionals brokering contact, we identified and contacted relevant organizations, explained our research plans, and asked relevant stakeholders to participate in an interview. In total, nine HDOs and three medical device manufacturers agreed to participate. During the interviews, we applied snowball sampling[94] by asking interviewees for references to colleagues in similar positions at the same or other organizations.

Notably, the roles and departments involved in software updating of medical devices varied across hospitals. Thus, participants with varying roles at HDOs and medical device manufacturers took part in our study. Participant demographics are depicted in Table 4.1.

4.3.2. STUDY DESIGN

To determine the research design, we began with an exploratory phase, in which we conducted pilot interviews based on our research questions and a review of previous literature. We interviewed six practitioners involved in security of connected medical devices, four of which were involved in the patching of medical devices at a hospital and two at a medical device manufacturer involved in product security with customer contact. During the interviews, we gained an initial understanding of the process of software updates for medical devices, stakeholders' responsibilities, and any challenges they faced. Based on these results, we designed the protocol for the final semi-structured interviews. Results from the pilot interviews were not included in the final analysis.

We designed a different interview protocol for HDOs and medical device manufacturers, as roles and responsibilities regarding security updates of medical devices differ between the two actors. However, as the focus of this work was to understand patching practices of medical devices at hospitals, we aimed at recruiting as many participants from HDOs as possible, while the manufacturer interviews provided additional context on the patching practices reported at HDOs.

After the first four interviews, we slightly adjusted some questions' wording and se-

Healthcare Delivery Organization participants:						
PID	HDO ID	Country	Role	Department	# cases	Experience
1	1	NL	Network Admin	ICT	-	Prefer not to say
2	1	NL	Biomedical Engineer	Medical Engineering	3	15 - 19 years
3	2	NL	Medical Physicist	Medical Engineering	-	10 - 14 years
4	2	NL	Security Officer	IT	-	1 - 4 years
5	1	NL	Biomedical Engineer	Medical Engineering	3	5 - 9 years
6	2	NL	Biomedical Engineer	Medical Engineering	1	15 - 19 years
7	3	ITA	IT/ICT Manager	ICT	-	15 - 19 years
8	3	ITA	Biomedical Engineer	Medical Engineering	-	< 1 year
9	3	ITA	System Admin	IT	-	≥ 20 years
10	4	NL	Medical Physicist	M-ICT department	-	≥ 20 years
11	4	NL	Biomedical Engineer	M-ICT department	3	15 - 19 years
12	5	NL	Biomedical Engineer	Medical Engineering	1	15 - 19 years
13	6	NL	Medical Physicist	Medical Physics	3	≥ 20 years
14	4	NL	Biomedical Engineer	M-ICT department	3	5 - 9 years
15	6	NL	Biomedical Engineer	Medical Technology	-	≥ 20 years
16	6	NL	Biomedical Engineer	Medical Technology	3	≥ 20 years
17	7	NL	Biomedical Engineer	Medical Technology	3	≥ 20 years
18	8	UK	Biomedical Engineer	Clinical Engineering	1	< 1 year
19	9	UK	Biomedical Engineer	Clinical Engineering	1	1 - 4 years
20	9	UK	Biomedical Engineer	Clinical Engineering	1	1 - 4 years
Medical Device Manufacturer participants:						
PID	Manu. ID	Country	Role	Department		Experience
21	1	DE	Product Security Specialist	Corporate Cybersecurity		10 - 14 years
22	2	NL	Product Security Specialist	Product Security		1 - 4 years
23	3	DE	Product Security Specialist	Product Management		1 - 4 years

Table 4.1: Participants' demographics. # cases indicates how many patch cases were mentioned during the interview. In case none were provided, the role did not implement software updates themselves. Experience refers to time in this role. For privacy reasons, role names were generalized.

quence but retained the overall structure and content. Slight adjustments of interview questions during data collection are accepted and even recommended for such action research by methodological literature (e.g., [154]), as it allows the researcher to respond to emerging, unexpected themes to probe for them more thoroughly in subsequent interviews.

4.3.3. INTERVIEW PROCEDURE

Opening the HDO interview after collecting informed consent (see Section 4.3.5), we began with general questions about the participant's role, followed by how connected medical devices and their security are managed within the HDO. We then asked interviewees to walk us through the process of how these devices receive security updates. At this point, we asked about details of the last three times they installed a security update on a medical device. Specifically, we were interested in the device type, how they learned about the patch, the timing of patch release and installation, the installation process, and how the device was actually connected to the rest of the HDO (e.g., to which networks and other devices).

This way, we captured a sample of 25 software update instances across eight different HDOs from 14 different manufacturers to ground stakeholder perspectives in practice. Twelve interviewees provided such update instances, but not all were dedicated security updates; there were also bundled software updates that might include security aspects, making it impossible for participants to entangle them. Participants not providing these update cases were either not directly involved in *installing* the update but more involved with the managerial decision-making or network surveillance or were not aware of any software updates on their fleet of medical devices. The column “# cases” in Table 4.1 denotes how many software update instances the participant contributed.

Interviews with participants from manufacturers followed a similar structure, beginning with the interviewee's role, followed by how frequently their products in the field receive (security) updates, how risk is assessed and decisions concerning security patches are made, how exactly updates reach devices in the field, and how they retain an overview of their products at HDOs. The full interview protocols for HDOs and medical device manufacturers can be found in Appendix C.

In total, three interviews (P2-H1, P5-H1, and P15-H6 & P16-H6) were done in-person and 14 remotely via video conference tools. Most interviews were done between one participant and the researcher, while during four interviews, there were between two (P3-H2 & P4-H2, P15-H6 & P16-H6, and P19-H9 & P20-H9) and three (P7-H3, P8-H3, & P9-H3) participants present due to participants' tight schedules or the necessity to involve colleagues from varying backgrounds. All interviews were conducted by the primary researcher in English, while one interview (P15-H6 & P16-H6) was conducted in Dutch with the support of a second dutch speaking researcher. Two participants (P1-H1 & P12-H5) could not participate in person or did not want to be recorded, so we asked the interview questions via email. Interviews took 51 minutes on average (Range: 28 - 65 min).

4.3.4. DATA ANALYSIS

We analyzed interviews using thematic analysis and continued data collection until reaching theoretical saturation in the HDO interviews when no new meaningful theoretical themes emerged from the data [99, 154]. Interviews were recorded and then transcribed. The transcripts were analyzed by the primary researcher via open coding, annotating any emerging themes and creating an initial codebook (with one coder regarded as being appropriate for this form of coding [32]). After the first five interviews, the code book was discussed with three other researchers with varying backgrounds and codes were adjusted to better fit the data. This process of validation and refinement of the code book was continued until theoretical saturation was reached. The final refined codebook was then applied to the previous interview transcripts to ensure standardized coding across all interviews. Notably, we did not reach full saturation with the manufacturer interviews due to the small sample ($n = 3$). We do also report these findings, as our goal was to (i) complement and provide commentary from a different perspective on our observations at HDOs and (ii) present novel insights for the academic community, as security research involving medical device manufacturers is very scarce. The final codebooks are provided in Appendix C.4.

4.3.5. ETHICS AND DATA PROTECTION

The study was approved by the researchers' institution's human research ethics committee. Before the interview, participants were informed via informed consent and orally by the primary researcher about the study's purpose, that participation was voluntary and not compensated, and the data collection process; Transcripts and quotes were anonymized and transcripts, audio recordings, and all interview responses were exclusively stored on a secured network at the researchers' institution. The paper draft was shared with participants before publication for review and potential retraction of any quotes or results. No participants requested any changes to quotes or results due to privacy and/or ethical concerns.

Four interviews were conducted with more than one participant present and thus at risk of negatively affecting participants' safety to speak up in front of colleagues. However, this multi-person interview format was actively suggested to us by the respective participants due to the distributed knowledge or time constraints. Thus, participants arranged and agreed internally to do the interview together as a group without the researchers requesting this. In each multi-person interview, every participant raised their voice, as they were experts on their respective topics.

4.4. RESULTS

To understand the process of patching connected medical devices at HDOs, we first examine the infrastructure the devices are embedded in, describe security update delivery pathways, and contextualize these findings with manufacturers' perspectives. When referring to specific participants, we use the denotation " $P_ -H/M_$ ", where $P_$ precedes the participant's ID and $H_$ or $M_$ the HDO or manufacturer ID, respectively.

HDO ID	Country	# Employees	# medical equipment	# connected med. equip.
1	NL	4,000 - 4,999	10,000 - 14,999	NA
2	NL	3,000 - 3,999	10,000 - 14,999	500 - 999
3	ITA	2,000 - 2,999	5,000 - 9,999	< 500
4	NL	3,000 - 3,999	10,000 - 14,999	3,000 - 3,999
5	NL	6,000 - 6,999	NA	< 500
6	NL	4,000 - 4,999	10,000 - 14,999	3,000 - 3,999
7	NL	4,000 - 4,999	15,000 - 19,999	3,000 - 3,999
8	UK	20,000 - 24,999	75,000 - 100,000	4,000 - 4,999
9	UK	25,000 - 29,999	75,000 - 100,000	3,000 - 3,999

Manu. ID	Country	# Employees
1	DE	60,000 - 69,999
2	NL	60,000 - 69,999
3	DE	10,000 - 19,999

Table 4.2: Overview of participating organizations' country of headquarters, size, and medical equipment inventory of HDOs. The numbers of medical equipment and connected medical equipment were estimations by participants

4.4.1. INFRASTRUCTURE

Due to substantial observed variance across HDOs in managing connected medical devices and their security updates, we start by describing HDO infrastructures, processes, and responsibilities to understand what is actually being patched.

DEVICE ESTATE OF PARTICIPATING ORGANIZATIONS

Table 4.2 depicts HDOs' sizes and reported inventory of medical equipment. HDOs participating in this research varied in size, ranging from below 3,000 employees to more than 25,000. This size was also represented in the number of medical equipment in use. We had no insight into exact asset management systems, but the numbers reported to us give a general indication of the magnitude of inventory size. At the HDOs in our sample, this was easily in the thousands, ranging from around 5,000 to more than 75,000 individual medical devices of varying sizes and uses.

When asked for the number of connected medical devices among the total number of medical devices, numbers diverged. For some participants, this was not actively tracked, or they were only able to answer for their own departments. Others, however, had a more precise overview, as connected devices were registered in their systems in some way (e.g., via asset management tools or MAC addresses in the network access control system). As a coarse estimation, the number of connected medical devices at the HDOs ranged from hundreds (e.g. H3), to thousands (e.g., H8). Participants generally expected these numbers to rise, as "(...) *it is very hard these days to buy equipment that measures something that is not connected to a network or a server.*" (P13-H6).

MANAGEMENT AND DISCOVERY OF DEVICES

All HDOs had an asset management system in place to track the inventory of connected medical devices. The majority used Enterprise Asset Management (EAM) software, in

which medical devices, spare parts, and maintenance tickets were tracked. Yet, localizing the physical devices could be non-trivial, even when listed in an inventory database. Participants from four HDOs (ID 4, 6, 8, & 9) mentioned that, especially for devices in high numbers, like infusion pumps or bedside monitoring equipment, it would take long, or almost be impossible, for technicians to locate all devices when installing updates, as they had no insight into the devices' (alternating) location, even when devices were visible on their network traffic. Two participants explained how they would love to put Air Tags on each device after the next update to be able to find them again.

No HDO reported actively tracking software versions, available updates, and vulnerabilities of all their connected medical devices, as they would for conventional IT systems. However, four HDOs (ID 4, 6, 8, & 9) reported tracking software versions of medical devices via their EAM software, which required manual input after each update. While participants generally voiced content with this system, they acknowledged that *"it's hard but people manage."* (P11-H4).

DEPARTMENT STRUCTURES AND INTERACTIONS

Traditionally, medical equipment at HDOs is maintained by medical technology (or "clinical" or "biomedical" engineering) departments. However, due to devices' increasing connectivity, IT expertise becomes essential. This merging of disciplines also became evident in our study. We thus report on the observed departmental structures and distributed stakeholder responsibilities concerning software updates of medical devices, as this is fundamentally different from organizational patch management of more conventional IT infrastructure, which is usually done by system administrators [26, 138, 225].

Several HDOs (ID 1, 3, 4, & 6) combined their IT departments with medical technology departments within the last years. This integration was either done by incorporating medical technology fully into the IT department, including them into an umbrella cluster with its own management, or extending the IT department's responsibilities to medical device security while keeping technical maintenance with the medical engineering department. The two UK HDOs were NHS Trusts and much larger organizations. Thus, they generally had more specialized technical and inventory management departments and cross-department committees and projects were appointed for specific tasks or projects to combine expertise from these different highly specialized branches.

Both IT and medical engineering shared responsibilities for connected medical devices, with the network infrastructure usually being managed by IT, and the physical devices and their interfaces to networks by medical engineering. However, they were commonly referred to as two different worlds with widely different approaches and expectations. This included mismatched processes (IT-based change management and SCRUM sprints clashing with medical engineering practices) or differing expectations for a medical device's lifespan (10 to 15 years seen as normal by medical technicians but long for IT) and patch frequency (IT departments favored regular updates like "Patch Tuesday," while technicians preferred fewer changes to functional systems). P2-H1 remarked; *"...the IT guys, they love updates (...), but us medical technicians, we think; 'Hey, this is working fine. Let's keep it that way!'"*

The installation of updates on medical devices was usually implemented by medical device service technicians, either from the HDO internally or externally from the supplier. Only on rare occasions was it mentioned that clinical users (i.e., nurses, doctors,

radiologists) would install any updates. This was even actively avoided in H4, where the central ICT department struggled to disable automatic update pop-ups on certain machines to stay in control of the installation timing. P11-H4 elaborated; “(...) *anybody who starts up the system and gets this message will be allowed to do the update, so we shut that off as much as possible.*”

The decision-making around updating was reported to be more distributed across stakeholders and departments. For regular updates, often the technicians decided independently if and when to install, or would oversee suppliers’ external technicians in their updating task. In case the supplier maintained the device, the HDO still had the final say in determining if and when to install updates. The HDO’s finance or procurement departments were involved in the decision of which maintenance service contracts to procure (which would determine update frequency and support duration), IT departments would sometimes detect vulnerable devices and demand patches to be installed, and three HDOs had a team of medical physicists responsible for device safety, who would support or oversee technicians in the decision to update. Importantly, the medical departments usually were reported to have the final say and could also demand a particular update. We go into more detail on the decision-making process around updating in subsection 4.4.3.

4

4.4.2. PATCHING PATHWAYS

To address our first research question (*How are connected medical devices patched within their operational environment at HDOs?*), we provide an overview of the observed different ways in which connected medical devices receive software updates. Based on the interviews with medical device manufacturers and HDOs, we identified four different pathways as to how an available software update, security-related or not, reaches medical devices in the field;

- (i) The manufacturer has a remote connection to the device and can make updates available in this way.
- (ii) Service technicians are sent out by the manufacturer or a certified service supplier, who install updates on-site.
- (iii) The HDO certifies their own technicians, who can, in turn, install updates, which are provided by the manufacturer or supplier.
- (iv) An exploitable vulnerability with a critical risk to patient safety (FDA: “*Uncontrolled risk*”) triggers a dedicated emergency update process, in which the patch is to be distributed within a shortened time frame (e.g., according to FDA, within 60 days[112]). This can be achieved by following processes (i), (ii), or (iii). We provide more details on this in Section 4.4.4.

In the following sections, we report on our observations pertaining to (ii), (iii), and (iv), as remote updating was barely reported as a common practice by participating HDOs.

CONSIDERATIONS FOR CHOOSING A PATHWAY

HDOs differed considerably in their approach to software updates for their medical equipment. Specifically, the level at which the organization managed and installed software

updates themselves varied, with some HDOs leaving the majority of updating activities to other parties, while others had their own processes and installed as many updates as possible in their own accordance.

We asked the participants from medical device manufacturers which pathway was most commonly chosen by their customer base. All three manufacturer participants reported that manual installation of software updates by service technicians from them or by third-party service suppliers were the clear majority. P23-M3 estimated this manual installation by service technicians across their products would constitute around 80%.

The ability for manufacturers to make updates available remotely to medical devices was reported to be substantially lower in demand, although being their preferred method, as it allowed them to better control software patching cadence, run remote maintenance and diagnostics, and save on the costly process of dispatching thousands of technicians for a software update. All three manufacturer participants also noted that with increasing size and financial resources, HDOs would be more likely to employ and certify their own technicians to install software patches.

Then, what did the HDOs in our sample say about their considerations as to which update process to choose?

INTERNAL MANAGEMENT – DRIVERS

HDOs 1, 4, 8, and 9 ($n = 4$) reported to prefer to have their own technicians install updates on their connected medical devices as much as possible.

A major consideration to keep maintenance and updating in-house was cost. Two HDOs reported it was most financially feasible in the long run instead of opting for a service contract with the supplier. Biomedical technicians usually had to be trained and certified for specific devices to be allowed to maintain the software. Thus, six participants, all technicians at HDOs, regularly attended training for different medical devices, which varied greatly by device (complexity), manufacturer, training scope, length, and required re-certification. Despite high certification costs, both HDOs reported that according to their calculations, maintaining the medical device themselves would still be cheaper.

Another driver to install updates in-house was control of the updating process. P2-H1, P6-H2, and P14-H4, who regularly installed updates on medical devices, stressed the importance for the hospital to control the timing of the update, as devices could not always be removed from service, which then required live monitoring of the device and patient. Instances of uncontrolled updates were given, such as a laptop as part of a medical device automatically initiating an OS update during operations or half a day of canceled patient appointments due to a medical user being prompted on the device UI to install an update, which, unexpected to the user, took several hours.

In fact, control over the installation process was one of the main reasons why HDOs in our sample actively decided against remote and/or automatic software update delivery to their connected medical devices.

EXTERNAL MANAGEMENT – DRIVERS

Four other HDOs (ID 2, 5, 6, & 7) reported that for most of their connected medical devices, they would outsource the software maintenance to the supplier via a service contract. This could be the manufacturer, a retailer, or a third-party service supplier.

The most commonly voiced reason for this was medical devices' complexity and safety risks. Four participants explained that for highly complex devices with a potentially difficult and/or dangerous maintenance process, such as those involving strong magnets or radiation, they would keep this process with highly specialized external technicians. This was confirmed from the manufacturer's perspective by P22-M2, who described that for complex devices like an MRI, extensive testing after the installation of an update is required to ensure the absence of any safety risks.

Costs were mentioned as a main consideration again: three participants (P3-H2, P13-H6, & P18-H8) mentioned expensive and re-occurring training and an economy of scale, where it would only be feasible to certify internal technicians if they would have at least a certain volume of devices. Two HDOs also mentioned the benefits of staying within an ecosystem of one manufacturer, which would include updates, but also bulk discounts, and a (relatively) simpler maintenance process.

4

4.4.3. PATCHING CONNECTED MEDICAL DEVICES

To further understand our first research question as well as our second one (*What kind of challenges do HDOs and medical device manufacturers encounter during this process and how are they mitigated?*), we adapt the five-step sequence of Li et al. for the updating process in organizational settings [138], which has been applied in other empirical work on patching practices (e.g., [65, 225]): learning about updates; deciding to update; preparing for installation; deploying updates, and; handling post-update issues.

Naturally, some of the steps differed considerably between HDOs who chose to internalize the updating process and the organizations who outsourced it (most notably in deciding to patch and patch deployment).

LEARNING ABOUT UPDATES

HDO participants reported widely different ways in which they learned about available software updates (including security updates) for their medical devices, especially when compared to conventional IT systems. Most HDOs reported to be in a somewhat passive position and would not proactively check for available updates on a regular basis for all devices. However, there were some differences in this process between HDOs who managed software updates in-house and the ones who outsourced this.

For HDOs managing the software update process for many devices themselves, technicians had to actively check for available updates, as notifications about updates, security-related or not, would not necessarily arrive as a dedicated message but often be posted on a manufacturer-specific web portal, where the technicians could then download it from. In this case, the responsible technician would have to manually check for available updates on the respective platform to decide if the update should be installed.

As most manufacturers were reported to have their own native platform environments and communication methods (e.g., by an updated document, or a dedicated post), this was reported to lead to a considerable burden for medical engineering departments, as the plethora of different platforms and documentations would render active update-tracking extremely time-consuming, approaching impossible, for the entire device base. P10-H4 elaborated; "(...) *we have about 400 different suppliers. I cannot go to all the platforms once a month to check if there's something new.*", with P14-H4 noting "(...) *it's mainly*

the management process it's not the installation process." For this reason, several technicians reported they would check for available updates more or less randomly, whenever they were able to find the time, as they were busy enough with other maintenance duties.

Two HDOs (ID 6 & 8) also reported that email notifications about device updates or emergent risks might at times not follow the defined process, but instead arrive with recipients who would not know what to do with the message, e.g., the medical user who initially purchased the product.

Regarding the actual method of notification, the most commonly mentioned medium across all HDOs was a direct notification (via email or letter) from the manufacturer or the service supplier ($n = 9$), followed by the active checking on manufacturer-native platforms ($n = 4$). Three HDOs (ID 1, 8, & 9) also reported on regular medical device security and/or safety alerts and reports from authorities, such as the Dutch healthcare sector CERT or the MHRA [157] and NHS [168] in the UK, which could refer to an available patch.

More ad-hoc update notification pathways were mentioned across interviews, although less consistently, including: from an external service engineer during their visit, during annual sales/service meetings with manufacturer representatives, or by network surveillance identifying vulnerable medical devices and inquiring about a patch at the manufacturer.

It was less likely to hear from participants that a device's interface would notify them of an update at start-up. Several participants mentioned Log4j in 2021, when involved departments had to actively enquire at manufacturers if their devices would be affected, and when they could expect a fix.

One HDO also reported they would not be notified about available software updates at all. They would install patches on IT infrastructure, but not on connected medical equipment, even though they voiced the desire to control this installed base more actively. *"We know that this seems incredible but with medical devices, manufacturers and suppliers is very difficult to talk to and ask them for [a new software update]"*.

DECIDING TO UPDATE

It was often not possible for participants to disentangle security patches from overall software updates, as they would be bundled together. As a result, the decision to update was often made based on non-security aspects such as features, bug fixes, or performance improvements, with security patches as an additional yet less visible part of the update. This was confirmed by all participants from manufacturers, who regarded this as common practice in the sector.

The decision-making of stakeholders at HDOs was thus influenced by other, non-security related factors. *"The (device), it doesn't get security patching, so Windows security patching at all, we leave it as it is. Only when it goes down or when something else is happening, then we think OK, now we can install some patches, but otherwise, we don't install patches on these systems."* (P2-H1).

Yet, several participants clarified that a dedicated security update without any changes to UI or functionality would just be installed. *"...if it's only (security) patching then uh, the normal process is executed, but if for example also the user interface is changed or settings may be changed, then we perform an additional risk analysis."* (P3-H2).

Due to the entanglement of updates with functionality, the effort to manage and decide on patching rose considerably for HDOs. Each update had to be assessed for effects on device behavior and interfaces and if a rollback process was in place. For this, considerable communication with other stakeholders was often required, such as asking suppliers for details, checking with medical device safety specialists, and discussing the changes with users; *“And then the medical physicists check the release notes. And if there are any user changes, we discuss it with the user. Like you get a new UI or you get a new button somewhere.”* (P11-H4). P10-H4 further explained that getting the necessary release notes for this decision was not always easy: *“We ask for release notes, which are always really difficult to obtain. And there’s a lot of companies that still say no release notes.”* This process was potentially multiplied for hundreds of different devices, as for all their software updates, a decision had to be made.

A further struggle and a reason to postpone updates was a complex and time-consuming installation process, from the actual installation to the communication needed with the medical departments. It was often preferred to wait for the next maintenance interval, while several medical technicians also represented the view: *“why disturb something that is working perfectly and create a risk that it is not working perfectly afterward?”* (P2-H1).

Costs played a substantial role here too. Some updates had to be paid for by HDOs. As updates or upgrades (e.g., to a newer OS) would be charged by device, three participants from different HDOs explained how these costs would escalate quickly in case the HDO deployed several devices, and thus being *“not a good business case”* (P3-H2). P14-H4 explained that he would often be asked by his colleagues from IT about why they do not update medical devices more often, to which he would reply *“...because if I install a KB from Microsoft, for the IT guys, just download it, install it, and done. For us, it’s paying a bill of 10,000 euros. And then installing it. And then we’re done. We have thousands of connected devices, each month, there are updates.”*

Importantly, security-related emergency patches with critical risk were reported to be free of charge, as regulations require manufacturers to respond to such risks if posing a threat to patient safety. However, if the device was running on an outdated OS without patch releases from the OS vendor, only a (paid) upgrade to a newer OS would ensure further security patches. Responsibility for patching is then intertwined: *“In the end, we are responsible to make the final decision that we say yes, it’s allowed to patch it, but in a sense... the essence is that the manufacturer or supplier decides this device has to be patched or updated; they have the lead in that process.”* (P3-H2).

PREPARING FOR INSTALLATION

The steps to prepare for installation of an update were mostly related to ensuring continued patient care. This required active communication, and sometimes negotiations, with the medical users of the devices and, if involved, external service technicians, as schedules were usually tight, with many devices running almost 24/7 to maintain patient care: *–“we have to go to every department, ICU the OR and try to arrange this 15 minutes and that takes to a lot of time.”* (P6-H2).

Patching was at times spontaneous and reactive (*“Sometimes we make an appointment with the nurses that they call when the patient is gone and then the guys come with*

the USB stick", P13-H6). Conversely, emergency patches installed at short notice could create friction, e.g., telling users; *"Sorry you can't use it now, and tomorrow it'll work. And of course they get mad but not much you can do about it at that time."* (P5-H1). A participant noted that the strain of negotiating with medical departments resulted in the IT department preferring to postpone updates to dedicated maintenance intervals.

The overall sentiment was that medical departments usually understood the need to update and would find a way to install the update, as *"...they know we have to do it."* (P5-H1). This went as far as scheduling downtime and not booking in patient care that uses a device, *"so (medical staff) know for instance on certain apparatuses they don't have to plan patients because then there is an update."* (P13-H6).

Testing updates before rollout is a common preparatory step for system administrators of conventional IT systems. [138, 225]. For connected medical devices however, HDOs usually did not test the updates before installing them, as this was done by the manufacturer and/or an external supplier. Therefore, several participants also mentioned they would always plan a roll-back process and back-up data before installing updates, to ensure that if the update would introduce any issues, one could go back to the previous version. Subsection 4.4.3 details on post-installation issues.

DEPLOYING UPDATES

As depicted in Table C.1 in the Appendix, update install frequency was generally reported to be low for connected medical devices, ranging from installing an update every couple of years (e.g., Device 3, 22), to approaching quarterly update cycles (e.g., Device 9, 18). While this corresponds to previous reports and warnings [37, 56, 115, 179], this frequency was reported to be due to costs of updates (be it an effortful and time-consuming process or the actual costs per update) and to the preferred approach to secure medical devices on the network level, which was also done to secure devices running on an outdated OS. *"Medical devices, we do not patch them regularly, most systems. No. No. So we try to have a different approach. We put them in isolated VLANs."*(P14-H4).

For all update installations reported, physical access to the medical device was needed. This could involve inserting a Flash Drive ($n = 10$), initiating the update via the device's UI ($n = 7$), or via an external Laptop ($n = 6$). Thus, internal and/or external technicians would usually have to go to each device individually to implement and/or oversee the update installation. In one instance for HDO9 (Device 25), it took six months until the external and internal technicians rolled out a firmware update to all 200 ECG devices.

Furthermore, several devices were reported to have dependencies to systems like servers ($n = 7$), which could lead to dependencies for update installation, such as that updates could be made available to individual devices via a central server, or that installation would have to be done jointly, which, in case of Device 13, would include an external technician coming in with the laptop and update, an internal technician overseeing the process, and a network administrator to take care of the server. Similarly, three HDOs (ID 7, 8, & 9) had a policy in place for some devices (e.g., infusion pumps) that all devices across the organization had to be on the same version to avoid confusion for users. Thus, all devices would have to be updated collectively. Section 4.4.3 details on some of the complexities relating to this process.

Where portable devices (e.g., ultrasounds, automated external defibrillators) could potentially be taken into a workshop, fixed-place devices with intense clinical up-time

(e.g., patient monitors, operating room ventilators), would require live (bed-site) monitoring by technicians and nurses during installation: “*We have a policy that we say we want to be at the place when it’s done, so we can explain things to the users.*” (P6-H2).

HANDLING POST-UPDATE ISSUES

P5-H1 reported he experienced post-installation issues two or three times during his time at the HDO, and in P14-H4’s experience, around 5% of software updates would introduce some sort of issue requiring rollback. Several other technicians never experienced this but heard anecdotes about it from their colleagues. Four of the update cases depicted in Table 3⁰ were reported to have led to post-update issues and had to be rolled back, for instance, due to unexpected incompatibility with existing hardware (Device 1), or users reporting unexpected device behavior afterward (Device 13, 16, 18).

Preparing for an update only to need to roll back the work was perceived as a very painful process, especially if “*We roll back the update until the manufacturer finds out the cause of the failed update and comes up with a solution.*” (P1-H1).

4

4.4.4. MANUFACTURER PERSPECTIVE

We conducted interviews with three product security experts at three large global medical device manufacturers to understand the software update processes across their product portfolio, the risk management and decision-making process prior to security patches, and their observations of what was fed back to them from their customer base. Thus, we could contextualize statements from HDOs and contrast perspectives between these two actors.

SECURITY PATCHING PROCESSES AT MANUFACTURERS

The different pathways of software update delivery for connected medical devices offered by manufacturers in our study are described in section 4.4.2. Participants explained that the release frequency depends on the modality and its OS but that generally, they aimed for regular update package releases (quarterly or bi-yearly), which would usually bundle validated OS, application, and library updates.

In case of an exploitable security vulnerability with patient-safety implications, all three manufacturers had a dedicated process in place to fix devices in the field within a defined timeline. In the medical device sector, such safety-related product warnings and changes after market entry are termed “*Field safety notifications*”, and participants estimated this would happen around once or twice per year across their entire product portfolio due to security-related risks.

The manufacturers determined internally if a vulnerability within a medical device and its underlying components required this dedicated patching process. To do so, a risk assessment process was in place at all three manufacturers, in which arising vulnerabilities would be regularly assessed for exploitability and patient safety implications by product security teams using the device’s software-bill-of-material (SBOM). Notably, security updates were not the only mitigation option, as the network connections and use context were also considered, which allowed for different mitigation approaches. For instance, P22-M2 explained how an MR residing in a secured room constituted a different security risk than a smaller, mobile, and easily accessible Ultrasound device. Thus,

mitigation could also take the form of mandated use, where the manufacturer would mandate the customers to ensure proper protection (e.g., restricting physical access, not connecting it to the Internet, or securing USB ports). Thus, the decision about vulnerabilities' criticality in any soft- or hardware component and the response were decided by the manufacturer on a case-by-case basis, with HDOs then deciding if and when to install any patches resulting from this risk assessment.

In case of regular updates, bundling of security patches with other updates was common practice, although P21-M1 clarified this would depend on the targeted software stack (i.e., dedicated security patches could be released for OS layers). As each change to the system had to ensure continued patient safety under the regulations, validating software updates in test environments was a crucial yet costly process. P21-M1 described a continuous internal balancing and negotiation act, where on the one hand, smaller and dedicated security packages were preferred from a product security perspective due to shorter installation downtime and customers not having to check for functionality changes (and thus, higher install rates), and on the other hand, the push towards reducing costs by validating separate updates together and dispatching technicians for only one bundle.

P21-M1 and P22-M2 stressed that software update delivery for their medical devices was challenging and costly, especially for emergency patches on a shorter time window, as thousands of technicians would have to be dispatched to customers worldwide with their own schedules and infrastructure, and that also for remote software updates, technically implementing a reliable pipeline was non-trivial.

CUSTOMER OBSERVATIONS

We were interested in how manufacturers observed their customers' practices, decisions, and expectations towards software updates and security. As explained in Section 4.4.2, all three manufacturers observed that patching via a remote connection was in substantially lower demand than manual update delivery via service technicians. In their experience, reasons for this were interrupted schedules when a user would install an update via the device's UI, a dislike for functional changes in updates, a lack of trust towards manufacturers to handle device or patient data appropriately, or the unwillingness to pay for remote updating, as this could also include deploying further infrastructure, such as gateway servers at the hospital site. This mostly overlapped with our findings from HDOs.

They also observed increasingly demanding security and privacy policies from customers the device would have to comply with and a rising demand for more regular and faster patching cycles, especially among customers' IT departments. At the same time, all manufacturers commonly observed their products at HDOs running outdated software and operating systems, thus not being fully patchable, as customers would opt against updating or upgrading the device. P22-M2 explained how instead, many customers would invest in network-based solutions such as micro-segmentation or privileged access management to secure devices, not trusting or not being aware of the security measures implemented into the device itself by the manufacturer.

REGULATORY RAMIFICATIONS

The majority of regulations for medical device security apply to medical device manufacturers. Thus, the development and maintenance of their products were driven substantially by regulations. While all three participants acknowledged their central role and responsibility for patient safety, they also voiced concerns regarding some regulatory consequences for them and the healthcare sector as a whole.

To comply with all relevant global regulations, their requirements would be mapped to derive a unified process. This required substantial legal and compliance resources and was reported to often lead to a lock-in, where there was limited freedom in design decisions. P22-M2 furthermore explained how this compliance focus could lead to a state where ensuring compliance could put pressure on the communication with customers, who might then fall out of contact with manufacturers and disregard their products' security implementations.

P23-M3 explained how increasing regulatory demands could significantly increase the costs for medical devices and thus healthcare as a whole, as it would require many manufacturers to completely re-design their products' architecture to establish a regular and faster patch delivery process while continuously validating updates. Several HDOs also shared such concerns about the rising costs of medical technology.

4

4.5. DISCUSSION

In this study, we explored current practices of how connected medical devices receive security patches from the perspective of their operators (HDOs) and manufacturers. We summarize our findings and derive suggestions for practitioners and for future work, yet are cautious to generalize our results or derive categorical recommendations, as this is exploratory work.

4.5.1. UPDATE PRACTICES FOR MEDICAL DEVICES

Addressing our first research question (*How are connected medical devices patched within their operational environment at HDOs?*), we identified four different main pathways; Via a remote connection to the manufacturer, manual installation by service engineers, either by the supplier or the HDO, or, in case of vulnerabilities with critical risk, via a field safety notification in a shortened time frame.

We found that HDOs in our sample rarely opted for remote delivery of software updates for their medical devices, even when available, thus requiring manual installation on a per-device level by service technicians. Participants from manufacturers verified this as the most common preference among their customers. This was mainly driven by HDOs' desire to actively control the update process, avoiding interrupted patient care, and concerns for unexpected functional changes in updates.

Instead of having connected medical devices on the latest software version, HDOs often preferred to reduce network exposure, usually by implementing network segments (VLANs) to isolate devices. This was also evidenced by low update install frequencies due to operational overhead and costs for regular updates, thus rendering network solutions more scalable. While this network approach adds an additional security layer and is generally recommended [46, 85, 135], there are also pitfalls. One vulnerable or weakly

configured device can compromise the entire segment, which requires active tracking of software versions, vulnerabilities, and configurations of all devices in the segment, which was reportedly not the case in our sample. Previous work investigating HDO networks [69, 135] reported on medical devices frequently residing in the same segment as other IoT or personal devices like IP cameras, printers, or smartphones, defeating the purpose of retaining segments for a single medical use case and increasing risk due to potentially vulnerable IoT devices.

Compared to prior work on patching practices of more conventional IT systems, we identified several key differences in connected medical devices that must be accounted for. Firstly, the stakeholder group of (biomedical) service engineers is responsible for maintaining medical equipment, which includes installing security updates. This is very different from the populations of system administrators and IT professionals studied in other works [26, 67, 100, 138], due to different technical backgrounds and perspectives on IT security. Secondly, while previous work has also identified coordination across actors in the healthcare sector to be crucial in rolling out patches to IT systems [66, 67], our work highlights how connected medical devices, in contrast to backbone IT infrastructure, are significantly more exposed to physical use, thus being potentially mobile and more difficult to access to install patches without interrupting patient care. Thus, patching of medical devices is less scalable and less doable remotely, these being two conditions usually met with conventional IT patching.

Finally, regulations and business practices of the medical device market differ substantially, a topic not broadly covered by previous empirical work. For instance, patching and update support are often part of paid maintenance service contracts, a typical practice in the medical device market, but less so for conventional IT patching. Furthermore, regulations have a stronger effect on patching medical devices than on traditional IT infrastructure, as manufacturers have to validate software updates to their devices to ensure continued safety, which can slow the process and increase costs. While recent work has studied the complex regulatory landscape of medical device security [141], we invite future work to study empirically how current and future regulations (e.g., the PATCH Act [2] or NIS2 [240]) interact with organizational practices of connected medical device security.

4.5.2. CHALLENGES

Previous work identified endpoint complexity [120] and a heterogeneous software updating process [100] for medical devices as major challenges for HDOs' security posture. By asking our second research question (*What kind of challenges do HDOs and medical device manufacturers encounter during this process and how are they mitigated?*), our results provide a more granular view on these challenges from the perspectives of HDOs and manufacturers.

Costs. Delivering a software update to medical devices in the field involves substantial costs. It can be an effortful and time-consuming process for HDOs due to non-standardized and heterogeneous notification methods, deciding on complex or incomplete information with mingled security and functional aspects, extensive preparations requiring active communication with medical departments, and a potentially effortful installation process that does not scale well across multiple devices. Furthermore, in-

stalling updates usually requires certifications or service contracts for HDOs, and manufacturers need to sustain substantial operational costs for maintenance via service technicians and engage in a continuous validation of updates to keep up with changes across medical devices' hard- and software components.

Managing infrastructure. Keeping an overview of and visibility into the installed base of connected medical devices, their software version, active connections, configurations, and available software updates can quickly overwhelm hospitals, especially if they are smaller and/or have limited resources to invest in tools and expertise to manage this. In our study, four of the nine HDOs reported they would attempt to keep up with the changing software versions of medical devices by renewing this in asset management systems after installation, but no HDO had a process in place to proactively and regularly check for available security updates across medical devices and vendors, as this would require substantial additional resources. Instead, even for the largest HDO in our sample, it was perceived as reactive “*firefighting*” (P19-H9) for the involved technical departments due to a lack of resources and strategic vision for proactive management.

Regulatory implications. While the global regulatory shift towards more extensive vulnerability management and security patching of connected medical devices (Section 4.2.3) is promising, our work suggests it poses significant practical challenges for manufacturers and HDOs. More frequent security updates can conflict with the operational reality of many devices, necessitating architectural changes for future devices to enable continuous update validation and delivery while maintaining patient care. Deploying field technicians for each update would not scale with the anticipated increase in patches and device volume. This shift could increase costs for manufacturers, particularly smaller players with less IT experience, potentially reducing market diversity. Newer devices will eventually incorporate such update-delivery pipelines for compliance, but existing ‘legacy’ devices in circulation will be used for decades, as HDOs (often with limited resources) might run devices without being willing or able to manage patching, given that HDOs have the final say in installing an update while regulations predominantly target manufacturers regarding product security.

Entrenched silos. Medical devices' growing connectivity increasingly merges medical engineering (safety, functionality, usability) and IT (network connectivity, software, security). We observed this coupling of different disciplines raising challenges for both HDOs and manufacturers, such as differing priorities and expectations, even if combined structurally as departments. As IT staff expected frequent patches and control over the medical devices like for any other IT system, clinical engineers were more reluctant to change a functional system. For manufacturers, different priorities between security and features can constitute a balancing act.

Several of our identified challenges are congruent with previous work on security updates from other domains like consumer IoT or organizational IT, such as the potential for update-induced bugs, unexpected functional changes, or difficulty in obtaining all necessary patch information, e.g., due to a lack of a centralized source or release notes [64, 77, 104, 121, 138, 225, 246]. We extend this work to connected medical devices, identifying challenges unique to this domain, such as difficulties in locating and accessing physical devices despite intense medical use and the need to define responsibilities across clinical engineering and IT stakeholders.

4.5.3. LIMITATIONS

Our research has several limitations. Firstly, the samples of HDOs, medical device manufacturers, and patch cases do not necessarily represent the general population. Participants noted country differences in medical device security and patching, suggesting future research opportunities. Furthermore, due to the smaller sample of manufacturers, we did not reach full theoretical saturation for these three interviews. Still, they are large, globally-active market players reaching thousands of HDOs. Thus, our results provide valuable insights into the centralized risk and vulnerability management processes and experiences with customers worldwide.

We relied on a convenience sample from a hard-to-reach expert population in a strained healthcare sector, which may have led to self-selection bias [116], as organizations more adept at the security of medical technology might have been more likely to participate. Furthermore, the sampled patching cases during the interviews were based on participants' retrospective memory, potentially affecting some claims. However, we structured those recollections to facilitate comparison between HDOs and devices, albeit based on estimates rather than verified numbers.

Lastly, four interviews included multiple participants (see Section 4.3.3), potentially introducing bias like groupthink or reluctance to speak freely. This format was suggested by participants however, signaling a willingness to share experiences in a group setting. During the interview, the interviewer also adhered closely to the protocol and directed questions to each participant to keep interviews comparable. Questions were answered directly by one participant, occasionally supplemented by another; this represented how answers to complex problems may require the expertise of multiple professionals.

4.5.4. RECOMMENDATIONS

Managing updates at HDOs: To reduce HDOs' burden to manage software updates, a structured process to evaluate available updates was reported to be essential. For instance, HDO4 defined a role (in this case, medical physicists) to evaluate or demand the necessary update details and had a check-list that was iterated over for each software update; this helped structure the decision and clarify which aspects of an update need to be discussed with whom. Channeling all communications with vendors regarding safety and security through one contact (e.g., one email address or one department), helped HDO8 to streamline communications. Manufacturers should support HDOs in navigating the complexities of varying notification channels and formats of available updates or vulnerabilities by standardizing them together with other vendors, which would reduce the burden for HDOs and thus lead to higher install rates. As seen in our results, SBOMs furthermore help manufacturers to render the complexity of a medical device and its components manageable [38] and deliver more specific alerts to customers, if need be.

Interdisciplinary collaboration: Due to the increasing entanglement of hard- and software-related responsibilities regarding connected medical devices, we also want to highlight to HDOs the need for close collaboration between departments. Initiating dedicated commissions, projects, or clusters between IT, medical engineering, governance, and/or procurement departments helped HDOs in our study to manage both hard- and software-related challenges and uncertainties by bundling the necessary expertise. Importantly, we found that merging departments did not necessarily result in more effec-

tive cooperation due to vastly different approaches, tools, and cultures. Instead, we recommend HDOs to acknowledge the different worlds of IT and medical engineering but establish structures that allow for mutual, regular cooperation and bilateral learning to improve the capabilities to manage the increasing connectivity of medical technology.

Non-disruptive update mechanisms: We found that a major reason for HDOs to avoid remote updating, which could enable regular security patches, was the loss of control over the installation process and potential disruptions to patient care, as well as potential distrust towards the manufacturer. Thus, providing remote update methods that retain control over the installation for HDOs could increase adoption rates and trust towards updates and manufacturers. Case studies [251] and guidelines [85, 86] provide insights into how a continuous patch delivery process can be deployed that allows for operators' flexibility to install. Yet, if costs for such a deployment exceed manual software installations via service technicians, HDOs will have less incentive to opt for it.

4

4.6. CONCLUSION

In this work, we studied organizational practices surrounding the security patching of connected medical devices in their operational environment at HDOs. We found that providing such systems with security updates is non-trivial for the involved actors due to challenges in tracking software attributes across a vast and heterogeneous inventory of medical devices, an increasing strain for technical departments at HDOs to manage this infrastructure, and practical difficulties in preparing for and actually installing updates amidst medical use, as this usually required physical access to each device.

While medical devices become increasingly connected and are exposed to an evolving threat landscape, new regulations push towards a more frequent and faster delivery of security patches for such systems. Our work highlights that patching comes at a cost however, as dispatching and/or certifying technicians to install updates did not scale well and that willingness among HDOs to adopt available remote updating capabilities for medical devices was low. It thus remains to be seen how the actors in the healthcare system will balance such regulatory requirements with upcoming costs and medical device security and, thus, patient safety.

5

CONCLUSION

This dissertation empirically assessed the legal concept of user expectations regarding IoT security by conducting quantitative and qualitative surveys with end-users, stakeholders at organizations, and IoT manufacturers. It presents three peer-reviewed studies in chapters 2, 3, and 4, with the goal to address the overarching research question:

What are users' expectations regarding preventive and reactive security measures of IoT devices?

This chapter concludes the dissertation by summarizing the empirical results from the three studies in relation to the main research question, reflecting on the practical and societal implications, and suggesting directions for future academic work to build on the findings of this work.

5.1. EMPIRICAL FINDINGS

CHAPTER 2: WHEN SECURITY AND PRIVACY FAIL: UNDERSTANDING CONSUMER EXPECTATIONS

In Chapter 2, we studied user expectations regarding reactive response to emerging security and privacy-related incidents with smart consumer products. To explore expectations in lieu of product liability, that is, when something 'goes wrong' with a product's security, we studied whether and how user expectations varied for different responses to such incidents by the device manufacturer and the user, as well as for different types of smart devices and incidents. By conducting a vignette survey on an online crowdsourcing platform with 862 participants from various regions (but predominantly Europe and the USA), we found that expectations regarding reactive measures to security and privacy incidents varied substantially depending on the context. Survey participants expected device manufacturers to respond in some way to emerging security problems, with patching as the most likely and highly preferred response. In contrast, for privacy-related problems, such as non-consented data flows, data suggested a learned helplessness, as manufacturers' responses were rated more unanimously, including not respond-

ing at all to data collection practices perceived as inappropriate. Different responses from the user also led to varying expectations. Demanding a refund for the device was considered the most appropriate response, while continuing to use it was deemed ill-advised for potential security risks, but more suitable for privacy risks, indicating how the threat model of the adversary (security: 'hackers', privacy: a company) affected the user's possible options when confronted with an incident. Finally, we found that the type of smart device and security and privacy threat led to subtle differences in expectations, with products like smart cars, security cameras, phones, or speakers eliciting more conservative, risk-averse expectations regarding the responses than for other devices, such as a smart washing machine.

CHAPTER 3: PREVENTING FAILURES: EXPECTATIONS OF SECURITY SUPPORT OVER DEVICE LIFETIMES

Chapter 3 pivoted from expectations regarding reactive security to incidents to expectations of preventing such incidents over smart devices' lifetimes. We positioned user expectations in the context of the European Cyber Resilience Act and conducted an online survey with 993 participants in five different European countries to measure them empirically, as the duration to prevent and mitigate security risks for IoT products has to correspond to the expected use times of the product according to the CRA. We found that reported use times of smart devices as well as expected lifetimes commonly exceeded the CRA's minimum support duration of five years, with lifetime expectations varying across device types. For instance, smart thermostats had a significantly longer expected lifespan than smartwatches. We also found that expectations are an inherently context-dependent construct, as not only device type led to different expectations, but also individual factors of the survey participants, whether devices were smart or non-smart, and the kind of factors people took into account when forming expectations. Surprisingly to us, a majority of respondents also expected that software update support would correspond with devices' full lifetimes already today, highlighting how the current market dynamic of short support times seemed to be in conflict with expectations.

CHAPTER 4: SECURING CONNECTED MEDICAL DEVICES: PATCHING PRACTICES AND EXPECTATIONS AT ORGANIZATIONS

In Chapter 4, we moved the scope to IoT devices used in an organizational context and the respective practices and expectations of professional stakeholders. We studied patching practices and expectations surrounding this process for connected medical devices in healthcare delivery organizations. We conducted 23 interviews with practitioners involved in the patching process at several hospitals in different European countries, further enriching the picture by involving the device manufacturers in the study. We identified key challenges in securing connected medical devices in their operational environment: managing a quickly increasing infrastructure of IoT systems put a significant burden on hospitals, as they often lacked the resources for thorough inventory and patching management, and installing software updates to individual devices often required substantial effort, as physical access to the device was needed in face of continuous medical operation and scheduling conflicts. Furthermore, we found that stakeholders' expectations regarding IoT device security played a role in shaping practices and business rela-

tions in the organizational setting. Different departments had varying expectations regarding security (e.g., patching) and medical device lifespans. While practitioners from a medical engineering background (e.g., medical technicians, biomedical engineers, or medical physicists) preferred a running system not to be changed (i.e., updated) and often saw more than 10 years as a reasonable lifetime, stakeholder from hospitals' IT departments usually expected shorter lifetimes and higher patching cadence for connected medical devices, as they were considered more like any other IT-endpoint. Interviewees from medical device manufacturers further highlighted how customers got increasingly demanding regarding security product requirements, and how distrust towards the manufacturer (expected privacy issues with patient data) and expected invasive automatic updates (interrupting operations and a lack of control) prevented them from opting for the more secure patch delivery option: remote and automatic patching.

5.2. DISCUSSION

We now reflect on the findings in Chapters 2, 3, and 4 with respect to our main research question: **What are users' expectations regarding preventive and reactive security measures of IoT devices?** We first consider reactive security and position our results within related academic work and ongoing developments in product liability legislation. We then turn to the observed expectations regarding preventive security and how they relate to ongoing discourses on maintaining it over devices' lifetimes.

AFTER THE FACT: USER EXPECTATIONS REGARDING SECURITY INCIDENT HANDLING AND REMEDIATION

When it comes to reactive security (i.e., *responding* to emerging vulnerabilities or actual incidents), our results suggest that users generally expect a response from the manufacturer or vendor of the IoT device in the form of a security update. In Chapter 2, we found that patching security vulnerabilities was the only manufacturer response that was perceived to be both likely *and* appropriate. A physical product recall, on the other hand, was deemed appropriate but unlikely for smart devices, and no response was considered both unlikely and inappropriate. We also found in Chapter 3 that the majority of participants expected that smart devices would receive updates over their entire lifespan, thus highlighting how this response was an integral part of expectations. In contrast, previous work has frequently reported on unpatched, vulnerable, and compromised IoT devices, thus indicating that vulnerability and incident responses like patching often do not meet such expectations in reality [153, 163, 176, 183, 197, 201]. Certainly, being able to expect a patch to fix a risky device would be a much more predictable path for users than taking matters into their own hands. We found in Chapter 2 that participants displayed uncertainty as to how a user should respond to arising security risks, with only not responding at all being perceived as clearly being risky. Previous work also showed how remediating (potential) security or privacy incidents remained tricky for users [27, 199], with some users being fine with continuing to use a device despite potential security or privacy risks [123, 262] or stopping using it altogether [27, 201, 224].

However, patching as a reactive security control is increasingly codified into mandatory legislation [2, 230, 243]. For instance, the CRA provisions that vulnerabilities should

be remediated without delay in relation to the risks to the device. Such requirements from laws like the CRA will also be considered for product liability after an incident has led to damage, as the Product Liability Directive now also considers them as part of "*the safety that a person is entitled to expect*", and manufacturers can be held liable for unpatched vulnerabilities that are "*necessary to maintain safety*" [242]. As recent research shows how liability is often excluded to the maximum extent in contractual terms by vendors of IoT products [143], this incorporation of security aspects into liability law seems necessary. Thus, our results suggest that these regulatory pushes towards more predictable reactive security measures seem to meet user expectations observed in this work as well as other academic research [101], where IoT device users stressed the role of the government to define standards and regulations for a level playing field.

5

Yet, it remains to be seen how effectively product legislation like the CRA, RED, or PLD can improve the state of reactive security behaviors in the IoT market. In Chapter 4, we studied patching as such a reactive behavior for connected medical devices. The medical device market has been regulated for much longer than other, more general-purpose IoT devices like smart home products considered in Chapter 2 and 3. In medical device regulation, patching is considered part of the continuous process to mitigate risks to patient safety [141, 232, 238]. Thus, our observations made in Chapter 4 can serve as a case study for how regulations targeting reactive security play out in practice. While stakeholders at HDOs generally expected that medical devices would receive patches in case of vulnerabilities, they also acknowledged their inability to keep track of all notifications and updates for their entire infrastructure and to patch all devices within a short timeframe due to the need for continuous uptime. Additionally, HDOs often wanted to maintain active control over the device's version and the updating process, and thus reported a low willingness to allow remote and automatic updates, even though it would increase patching cadence and vulnerability remediation rates. The common practice of bundling security with features into the same update further complicated matters for HDOs, as feature changes had to be scrutinized before installation, and if undesired, led to not installing the latest security updates.

These observed challenges in implementing the required reactive measures will likely also apply to other, non-critical IoT products covered under more recent regulations, as they highlight how users' choice and preference can affect reactive security behaviors, rather than legal provisions or technology. The number of smart devices also rises in consumer households, requiring more active "inventory management" by users, and some users dislike updates to their devices [102, 246]. Thus, a certain set of users will not be able or willing to keep track of their IoT devices' updates, or switch automatic updates off to stay in control of the device's functionality and maintain the paradigm "never change a running system". Thus, Chapter 4 demonstrates that despite regulatory oversight and resulting high expectations as well as the efforts of manufacturers and customers (here, hospitals), significant challenges remain for remediating (i.e., patching) vulnerabilities in IoT devices deployed in the field.

PREVENTATIVE SECURITY: USER EXPECTATIONS REGARDING SECURITY OVER IoT DEVICE LIFETIMES

To prevent incidents for IoT devices, security controls must be upheld after the product has entered the market to manage and address arising risks. Defining an appropriate duration for this security support and how to actually uphold it is an ongoing discourse in academic research, industry, and policy. For instance, Bradley et al discuss how vendor dependency for IoT devices could be reduced for longer support durations [30], IoT vendors more frequently seek business models for post-sale revenue, which can finance ongoing costs of security support (e.g., via paid subscriptions or advertisements [108, 213]), and policymakers draft laws for how to support IoT devices over their lifespans. In the EU, the CRA defines support durations for different product categories, and also in the USA, there are ongoing policy initiatives for longer support times for consumer IoT devices [33, 209]. Both the CRA and model legislation in the US tie support durations to reasonable user expectations, thus bringing attention to consumers' perspectives and this legal concept in general.

Contributing to these ongoing developments, we studied user expectations regarding support duration in Chapter 3 and highlighted several characteristics of these expectations. We found that expectations varied depending on the type of device, with more consumer-level smart devices, such as smartwatches or smart speakers, displaying significantly shorter expected lifespans and usage times than more white-goods-type smart devices, like smart washing machines, thermostats, or WiFi solar inverters. Contributing to these device differences, we found that users took different factors into account for different device types (Table B.2 in the Appendix), showing how (expected) usage patterns and built characteristics led to varying appraisals by users. Such differences highlight the challenge in defining a single, horizontal minimum support duration applicable to all types of IoT devices and the importance of user expectations as a flexible concept for capturing variance across IoT devices. In a similar vein, we found that almost all IoT devices considered in the survey were expected to last longer than 5 years on average, thus exceeding the CRA's horizontal baseline of 5 years.

If the IoT market should gravitate towards this five-year period as the de facto standard support duration after the CRA comes into force, our results suggest that user expectations will not be met for a wide range of product types. Furthermore, as we found that actual use times also frequently exceeded five years, usage will likely resume after the end of support for a sizeable number of smart devices, thus maintaining the state of unsupported and vulnerable IoT devices. Interestingly, while short or inconsistent support times have been common in the IoT market [81], we also found in Chapter 3 that the majority of participants expected smart devices to be supported with updates for their entire lifespan. This expectation may also be reflected in strong negative public reactions towards IoT vendors stopping support or "de-smarting" their products at some point [18, 106, 107]. Vendors in such examples frequently cite incompatibility of the device with more recent technology or the upkeep costs of the cloud infrastructure as reasons to stop support.

This highlights a key challenge in IoT security: continuously maintaining and securing IoT devices does not have a strong business case for manufacturers. While legislation like the CRA attempts to enforce it by law, the case for the continuous mainte-

nance costs and effort persists. In Chapter 4, we also observed similar challenges for IoT medical devices that have already been regulated for many years. While expectations between stakeholders in the organization regarding device use times differed (e.g., between biomedical engineers and hospital IT departments), connected medical devices were frequently used for substantial periods of time, sometimes without being patched, even if a patch was available. This was also validated by medical device manufacturers, who reported that long use times past their end of support and an unwillingness to pay for a patch were common behaviors among their customers. Apart from varying customer expectations and device use times, manufacturers also highlighted the economic costs of securing their products over their expected lifespan and how these costs had to be passed on to customers (and thus, to the healthcare system in general).

Thus, looking ahead at a more regulated general IoT market, it will remain an ongoing challenge to balance long device use times and the respective user expectations with the continuous effort, care [125], and costs that security requires for IoT manufacturers. As legislation will require prolonged support, the resulting costs will be passed on to consumers, assuming the manufacturer actually invests the necessary resources to comply with the law. While previous work suggests that users are indeed willing to spend more on secure IoT products [72, 95, 250], it remains to be seen how spending on IoT devices will develop under stricter regulatory oversight and to what extent IoT devices will be used past their end of support.

5

5.3. GOVERNANCE IMPLICATIONS

Building on the empirical insights gained from the three studies presented in this dissertation, we now reflect on possible governance implications for key actors in the IoT market: governmental bodies, IoT device manufacturers, consumers, and healthcare delivery organizations. Governance refers to the collective activities and efforts of different entities to solve societal issues or create societal opportunities, which not only include governments and respective administrative authorities, but also other public organizations, the private sector, and civil society [155]. According to Meuleman, interactions between these different entities can be categorized into three different styles of governance: hierarchical, network, and market governance [155]. Hierarchical governance is characterized by a top-down approach, where control is established through authority, rules, and formal procedures. A government regulating a market with laws and enforcing them is one instance of this. Network governance stresses coordination, trust, and shared values among independent actors, where power is distributed and cooperation thus becomes crucial. Market governance is driven by the core idea of free markets, where economic dynamics drive actions and cooperation, specifically demand, supply, competition, and self-interest. Actors make decisions based on costs, incentives, and contracts, where the market is assumed to allocate resources efficiently.

GOVERNMENTAL BODIES

The IoT market has so far largely operated under a market governance paradigm, where competition among a plethora of IoT device vendors has led to a prioritization of fast time-to-market and new features, with less focus on security. As we have seen through-

out the different chapters, this free market approach has led to externalities in the form of vulnerable IoT devices and societal risks, triggering policymakers to intervene. We are currently in a transitional phase in the IoT market, where the market is increasingly shifting from market governance towards hierarchical governance. New legislation, such as the RED [230], the CRA [243], the UK PSTI act [229], or the delegated Product Liability Directive [242] are prime examples of how governments now exert authoritative power by introducing compulsory security requirements and rules for security-related liability within legislation. Previous work suggests that such hard pushes towards mandatory compliance can indeed lead to improved security outcomes, such as more thorough incorporation of security during product development [145, 227, 255], or an improved availability of IoT products' security disclosure on online retailers [249]. In Chapter 2 and 3, we also observed that several aspects of the upcoming law's product requirements appear to align with users' expectations. Manufacturers were expected to patch arising security vulnerabilities in their devices (Chapter 2), which corresponds to the numerous stricter patching regimes outlined in legislation, such as the CRA, RED, UK PSTI, or PLD. Many users expected long IoT device lifespans, used their devices for many years, and expected software update support to last for device's lifetime (Chapter 3), which meets the CRA's focus on continuous risk mitigation throughout product use times, but is in stark contrast to the current (inconsistent) market practices when it comes to update support [81].

Nonetheless, several open questions and gaps remain between our empirical measures and legislation, with implications for governmental bodies such as market surveillance authorities or policymakers. In Chapter 3, we saw that the five-year baseline for security support, as mandated by the CRA, was insufficient for the vast majority of IoT devices' expected lifespans, with planned use times also commonly exceeding this five-year threshold. Thus, taking these findings into account, the risk persists that millions (or even billions) of IoT devices could re-enter the state of being unpatched and vulnerable after the end of the minimum support duration, if people continue using them. We also found in Chapter 3 that the end of update support was only a weak driver for users to stop usage. The CRA leaves the room for the European Commission and the ADCO to release guidance on the minimum support durations for different product categories in the future, or even to delegate the Act. Thus, our work suggests that policymakers and market surveillance authorities involved should consider empirical data on IoT device usage times (expectations) when enforcing the law and publishing guidance. This will reduce potential gaps between actual use times and shorter support times for IoT devices, thereby satisfying users' expectations while mitigating societal risks associated with vulnerable IoT devices.

Ultimately, implementing and enforcing this new legislation to govern IoT security in practice will be a significant challenge in the years to come. In Chapter 4, we found that despite heavy regulation in place, it was still challenging for the involved stakeholders to secure connected medical devices in practice. For instance, the costs and added procedures of validating and verifying patches to comply with the Medical Device Regulation [232] led to patching delays and significant costs for manufacturers, with the theme emerging that smaller manufacturers might struggle to design and maintain connected medical devices that meet the strict requirements. Additionally, these rising costs

of security compliance are passed on to the customer (in this case, the hospital), thereby increasing costs for the healthcare sector as a whole. As regulations like the CRA or the Cyber Security Act [235] will require security certification for some IoT devices with certain risk levels, rising costs of compliance are likely to also develop in the general IoT market. With previous work discovering various vulnerabilities in *certified* IoT products [144], it remains an open question how market surveillance authorities will oversee the market for IoT devices and certification in practice. While only the courts can determine what constitutes a reasonable expectation for an IoT device's security, our results can at least provide market surveillance authorities with a relatively independent indication of the abstract notion of users' expectations when enforcing the law without having to resort to court. For instance, when discussing with manufacturers what constitutes a reasonable support duration, or what responses to security vulnerabilities are appropriate and expected.

IOT MANUFACTURERS

5

For manufacturers of IoT products, the regulatory landscape that need to be navigated is changing rapidly. While previously, mostly forms of market and network governance applied to the market IoT manufacturers, hierarchical governance in the form of regulations now requires adaptation. In the EU, the Cyber Resilience Act in particular will significantly affect the IoT market, as mandatory requirements will require manufacturers to consider security risks throughout the product's expected lifespan.

This work has several implications for manufacturers of IoT devices navigating these upcoming hierarchical governance forms. Users' expectations as a concept are relevant in both product conformity and liability legislation and thus need to be considered by manufacturers in some capacity. When it comes to security support durations under the CRA, manufacturers have to document how user expectations are incorporated in the determination of the support duration. As such, Chapters 2 and 3 provide insights into user expectations, as well as methodologies for measuring them empirically. For instance, we saw in Chapter 2 the potential differences between normative and likelihood expectations, highlighting the importance of clearly differentiating between the two. Additionally, Chapter 3 presents a methodology in the form of a survey that provides quantifiable insights into user expectations of IoT device lifetimes, which can be extended by manufacturers to other categories of IoT products beyond those included in the survey. As market surveillance authorities can inquire about documentation regarding how user expectations were considered during the determination of the support period, this can provide a useful approach for manufacturers, especially before official guidance on support durations from the European Commission is available.

For instance, manufacturers could leverage their customer-facing or user-facing departments, such as sales or market and user research, to initiate a conversation with their users and determine and document their expectations. However, expectations can differ widely across individuals and the way they are elicited, as highlighted by all three studies. Thus, if a manufacturer were to follow similar approaches to measuring users' expectations as we did, they should account for variability in responses by considering sufficiently large sample sizes and effective segmentation strategies, such as gathering data points from diverse countries and user groups. Yet, user expectations are only one

of the criteria for determining the support duration under the CRA: the nature of the product and other relevant Union law are the other two mandatory criteria. This work then also provides empirical information for manufacturers regarding the *product nature* criterion, as we illustrate how different IoT product types varied according to their nature (e.g., regarding how manufacturers should respond to security events for different product categories in Chapter 2, or how long their expected lifespans are in Chapter 3). Chapter 3 furthermore suggests that the brand and the associated trustworthiness of the manufacturer were a major consideration for consumers to navigate security risks (as depicted in Figure 3.8), thus, highlighting the commercial potential in building trust with users.

Our work also has implications for how manufacturers can develop mechanisms to keep users informed about the security of their devices. In Chapters 2 and 3, we saw that (in-)security was often opaque and imperceptible for users. Chapter 2 highlights the imperceptibility of potential security incidents or problematic data flows, and Chapter 3 shows that for many IoT devices, users were unaware whether their own devices still received software updates or if they had actually experienced a situation where their devices were no longer supported with security updates and risk monitoring. The CRA targets these challenges with extensive provisions on how manufacturers should disclose a product's security properties to users in Annex II. Thus, the three studies in this dissertation further stress the importance for manufacturers to take this seriously and make notification channels to users regarding potential security risks and available updates perceivable and actionable. Previous work reports on how security and usability departments collaborated to design privacy settings in lieu of the GDPR [43], which could now also be applied to security configurations and notification channels in the context of the CRA.

Notifications to users are especially crucial at the time when the support duration (i.e., the supply of security patches) stops. Once this point is reached, users can decide for themselves whether they want to continue using an unsupported device. We found that this was often the case. The end of update support was a weak trigger for consumers to stop using a device (Chapter 3), and in Chapter 4, we observed that this was also the case in an organizational setting, where users partly opted to run unsupported medical devices or not update them. This was due to unwanted functional changes introduced by an update or other, more scalable risk mitigation mechanisms in place, such as network segmentation. Thus, while organizations indeed have more technical tools at their disposal to mitigate security risks once an IoT device reaches the end of support, it is especially crucial for manufacturers of consumer products to provide clear notifications about the end of support to their users, so that they can make an informed decision whether to keep using it. If support times indeed increase due to the CRA's effect and end-of-support notifications become more noticeable, it will likely serve as a stronger trigger for consumers to discontinue device usage. Furthermore, the observed reluctance among some healthcare delivery organizations towards patching due to functional changes to the device is also reportedly present among consumers [102, 150, 246]. This highlights the potential security benefits of manufacturers releasing security updates separately from functional updates, as it can increase installation rates, both within organizations and for consumers.

CONSUMERS

We now turn to the direct and indirect implications for users of consumer IoT products. In the consumer IoT market, market governance has been the predominant governance form, which has been characterized by a lack of regulatory oversight. Thus, users largely had to organize in a manner akin to network governance: exchanging information and resources in a somewhat loose, non-standardized manner to reach common goals. Governments provided general guidance to users, while manufacturers provided varying levels of information regarding the security of their IoT products. Users often relied on community-building efforts to help each other, such as online forums [111, 201, 211] or seeking advice from friends and relatives [57, 187]. Arguably, this form of network-based interaction between actors will remain a significant form of governance even after the introduction of more thorough hierarchical governance. However, the enforcement of new laws will alter actors' incentives and the way interactions unfold, as well as who will bear responsibility for what. This work offers insights into the previous interactions between users and manufacturers, as well as their expectations regarding IoT security and responsibilities.

5

The primary focus of this research was user expectations regarding the security of IoT devices, which serve as a means for the law to establish a baseline for a given context, referring to society at large. Our work highlights several key details to consider when using user expectations as a baseline in the context of IoT security. First, expectations are commonly based on prior experience, that is, on the previous market dynamics of IoT security. This dissertation highlights the risks associated with relying on this baseline, as it can be self-defeating for users. If a previous problem in the IoT market has been persistent enough for users to become accustomed to, it shapes their expectations. For instance, in Chapter 2, we found a certain learned helplessness among users regarding security and privacy issues, manifesting in a somewhat cynical mindset among some participants, and the expectation that the manufacturer wouldn't do anything about problems, because they previously haven't done so. Similarly, Chapter 3 demonstrates how users factor in planned obsolescence and end-of-life through incompatibility when forming their expectations of products' lifespans, as they have experienced these issues in the past. For these reasons, smart devices had a lower expected lifespan than conventional, non-smart counterparts. These findings highlight how expectations based on what might be realistic to expect can actually not be in the own interest of users and society at large, as perverse market dynamics have shaped these expectations.

Furthermore, this work highlights the heterogeneity and volatility of user expectations. They are highly dependent on the way survey questions are phrased (e.g., shown in Chapter 2) and exhibit high variability across different individuals (Chapter 3) or departments in organizations (Chapter 4). It is therefore crucial to note that users can be nudged towards a certain (desired) answer when being surveyed about their expectations, and that these expectations can take different shapes depending on who is being queried. Thus, expectations as an empirical construct cannot, and are not, the sole criterion used as a benchmark in product conformity and liability law. For instance, the CRA explicitly refers to the nature of the product and other relevant union law as other mandatory criteria for determining the support duration. In Chapters 2 and 3, we indeed observed substantial differences in the data between product categories, due to their

varying "nature": different IoT devices elicited different responses regarding emerging security and privacy risks (Chapter 2) as well as expected and actual use times (Chapter 3). Thus, expectations can be shaped by a low bar of security in the past IoT market and can be manipulated depending on how they are measured. Such risks of empirically measuring user expectations in law have also been raised by legal scholars [91] and constitute a key societal implication of this dissertation. Thus, when user expectations are deployed as a baseline by manufacturers, market surveillance authorities, and eventually, by the courts, it will be crucial to consider how they have been determined and that they are shaped by past market conditions that might not be in society's best interest. Thus, also introducing a certain normative component when conceptualizing expectations is crucial: reasonableness might not only be based on how things have played out in the past, but also on what people *should* be entitled to expect.

This work also has implications for how consumers interact with the legal framework more generally. With the upcoming regulations, consumers of IoT products in the EU will have more legal rights. For instance, users can have an easier path to proving security-related damages in their IoT device for liability [242] or demand information from the IoT manufacturer about the security properties of the product [243]. However, our work raises the question of the extent to which users will become aware of these additional rights, recognize emerging security issues or defects with the products they are using, and whether they can and will actually pursue legal action. In Chapter 2, we observed learned helplessness among users regarding the remediation of security or privacy problems with their smart devices, and even instances where users did not know any other option to respond to a compromised device than to dispose of it. Thus, it will be crucial to ensure that users have predictable processes for asserting their rights, and if these are not met, to pursue legal action. Arguably, if such mechanisms exist for a long enough period, awareness of them and expectations for them will likely also develop, and users' sentiments of helplessness will decrease. Of course, actually implementing actionable legal paths for users in practice will ultimately depend on the effectiveness of the respective judicial systems that enforce and interpret the law.

HEALTHCARE DELIVERY ORGANIZATIONS

The digitization of medical equipment is progressing rapidly in hospitals, and securing this infrastructure will be a key governance challenge in the coming years. In the medical device market, all forms of governance [155] can be observed: market dynamics in the form of competition across medical device manufacturers, network governance in interdependent interactions between HDOs, manufacturers, patients, and governmental and public agencies, as well as hierarchical governance forms, as the sector is heavily regulated [141, 232].

In this complex, multi-actor system, collaboration between stakeholders and actors is crucial in reaching better security outcomes, despite regulations provisioning rules on how security ought to be handled in theory. In Chapter 4, we found that medical devices can only be reliably and continuously secured (i.e., patched) if product security experts at manufacturers, external or internal service technicians, HDO IT and medical engineering departments, and the medical units work together in a highly collaborative and interactive manner. With different expectations, approaches, and priorities, these differ-

ent stakeholder groups are needed to manage the technologically and logistically complex task of managing and securing the infrastructure of connected medical devices and their operational (network) environment. We found that, despite regulations in place, this process was often difficult and very effortful to implement. Key challenges observed included an increasingly complex inventory and patch management that overwhelmed hospitals, notification channels between actors that were frequently unreliable, costly certification and verification processes, and patching rates that barely reached 100% across manufacturers' customers.

Thus, our results imply that HDO's needs and perspectives have to be considered more thoroughly in finding patching processes that are workable in practice. It has been reported to us that the most common method for delivering security updates to medical devices is through the dispatch of service technicians to install updates on individual devices. With the plethora of different soft- and hardware components integrated into a modern medical device and a generally rising volume of vulnerabilities, this method does not scale and will likely have to be replaced with a faster and more reliable way of patching, such as over-the-air updates. However, we found that most hospitals were opposing this due to first-hand experiences with automatic and remote updates introducing unwanted system changes, interrupting patient care, or concerns regarding remote access and patient data protection. Thus, patching methods will have to be developed in the healthcare sector that allow for the unique requirements of patient care at HDOs: giving HDOs control over update timing, providing clear and standardized information on all changes to the device enclosed in an update, and debundling functional from security and bug-fixing patches as much as possible. This will require close collaboration and negotiation within medical device manufacturers (e.g., between security and product teams) but also between manufacturers, HDOs, and regulators, to meet HDOs' (i.e., the users') operational needs and to balance regulatory requirements for repeated certification and validation with the volatile nature of software and security vulnerabilities.

5

5.4. FUTURE WORK

This work has empirically studied user expectations of IoT security in the context of substantial regulatory developments. As the rollout of regulations for all kinds of IoT products is just beginning, many open questions remain about how the law can be applied in practice and what effects it will have on the IoT market. Thus, future research is needed to support the necessary governance mechanisms that will apply the law. This includes a more thorough empirical investigation of legal concepts (such as user expectations), as well as the behaviors of manufacturers, market surveillance authorities, and users in the context of this new legislation.

A CLOSER LOOK AT USER EXPECTATIONS OF SECURITY

While this work has investigated user expectations of IoT security, it remains mostly exploratory in nature. Thus, future work could study expectations more deeply and consider more factors contributing to this highly context-dependent construct. For instance, liability law considers various factors surrounding a potential product defect to establish what is a reasonable expectation regarding the safety of this product. This in-

cludes the device's price, marketing presentation, other comparable products, or how it was used. The causal relationships between these factors and user expectations in the context of IoT security could be further established experimentally through vignette experiments similar to the one described in Chapter 2. That is, what would be the effect of an IoT device's price or "premium-quality" on expectations about its security? As we found in both Chapter 2 and 3, brand was a factor considered by users when it comes to security, but which was not further studied in detail.

Furthermore, there remains potential in researching user expectations in an organizational context. While we looked at Healthcare Delivery Organizations and Medical Device Manufacturers in Chapter 4, future work could extend this line of research to organizations deploying an IoT-driven critical infrastructure in other domains, such as industrial manufacturing, the logistics sector, or energy grid providers. Stakeholders' expectations regarding their devices' security features, patching regimes, compatibility requirements, and lifetimes could prove to be interesting contrasts to more consumer or medical IoT expectations, due to different contractual and operational conditions. Such research would furthermore support ongoing policy discussions regarding the support durations of the CRA, as it would empirically establish (lifetime) expectations for such specialized IoT products, if they fall under the CRA.

EMPIRICALLY ASSESSING OTHER ABSTRACT LEGAL NOTIONS

Future empirical research can continue to support policymakers, market surveillance authorities, and IoT manufacturers in alleviating some of the uncertainty surrounding legal concepts that are crucial for applying the law but remain largely abstract for practice. While not legally binding, such input from independent research can support ongoing discussions among these actors regarding workable solutions for applying and complying with the law in practice. For instance, beyond reasonable user expectations, the CRA emphasizes the concepts of *reasonably foreseeable use and misuse* that manufacturers must consider during the risk assessment of their products. In the definition of both use and misuse, the CRA mentions that this refers to behavior that is "*likely to result from reasonably foreseeable human behavior or technical operations or interactions*". Empirically exploring the lines between foreseeable use and misuse by assessing how IoT devices and security features are used and appropriated by users in practice could be another promising research direction that would directly support manufacturers in their risk assessment under CRA compliance, as it could provide real-world instances as illustrative examples for practitioners during risk assessment.

Open questions also surround the concept of *substantial modification* within the Product Liability Directive [242] and the CRA [243]. When a product is "substantially modified", the expiry period of liability starts from zero again. For IoT devices, this raises the question of when a software update can be considered a substantial modification to a product, as the PLD explicitly states that "updates or upgrades that do not amount to a substantial modification of the product should not affect the expiry period", implying that updates or upgrades *can* indeed amount to a substantial modification. The CRA includes this concept too, as a substantial modification would require re-certification and a restart of the support period counter. While the CRA goes into more detail on the definition and applicability of substantial modifications in the context of (security)

updates¹, it would be interesting to see when users perceive an update as a substantial modification, and if and how such substantial functional changes to a product affect updating behavior or willingness. This could be an interesting research direction for both empirical and legal researchers, as it directly connects the law with real-world instances and can thus support its implementation. For instance, the CRA states that the European Commission will also publish guidance for manufacturers on this very concept (Article 26, 2 (d)), which has not yet occurred.

ACTORS' BEHAVIORS AFTER REGULATIONS HAVE COME INTO FORCE

Besides informing ongoing policy discussions surrounding legal concepts, future work should also empirically study the effects of recent regulations on the IoT market. Thus, it can measure the law's effectiveness in terms of security and market outcomes, and inform policymakers about potential pivots, such as releasing specific guidance or revising a law.

Potential research following up on this work could investigate how IoT device certification under the RED and CRA will play out in practice, as we observed in Chapter 4 that security challenges remained despite extensive certification and regulation in the medical device sector. This includes research questions regarding the availability, effectiveness, and running costs of notified bodies conducting device testing, and to what extent security will improve for certified IoT products. Recent research questioned the impact of certification on security [144]. Related to this performance of the certification sector, research is also needed to support market surveillance authorities in market oversight and law enforcement. With millions of different IoT device models being introduced to the European markets, processes and methodologies are required for market oversight agencies to handle this deluge of products. To make things even more complex, most of these IoT products are likely to undergo continuous modification through software updates. Research can support this endeavor by introducing empirical-based and risk-based approaches on how to conduct oversight for such a heterogeneous and complex market, for instance, by developing scalable scanning methodologies of IoT devices and their security attributes.

Furthermore, researching consumer behavior and expectations after the new regulations have been in effect for some time will be relevant in understanding the effectiveness of the regulations and their impact on the public's perceptions of security. Security and risk awareness may change among the population as security regulations become increasingly stricter. For instance, it will be interesting to see if users would be more likely to stop using their IoT devices after the end of support than we have observed in our studies, as notifications about this will become mandatory and hopefully clearer for users to notice. This question would also be interesting to study in terms of the intended purpose of the IoT device at hand. For an IoT device where smart features are not integral to the intended purpose but a "nice to have", the end of security support would constitute an entirely different situation than for products that require an active internet connection to function as intended. A smart washing machine may still function properly after being disconnected, unlike a smart speaker that usually requires a connection for the key features.

¹Specifically in recitals (39) up to and including (42).

BIBLIOGRAPHY

- [1] 116th Congress (2019-2020). 2020. H.R.1668 - IoT Cybersecurity Improvement Act of 2020. <https://www.congress.gov/bill/116th-congress/house-bill/1668>
- [2] 117th Congress (2021-2022). 2022. S.3983 - PATCH Act. <https://www.congress.gov/bill/117th-congress/senate-bill/3983>
- [3] Denielle Abaquita, Paritosh Bahirat, Karla A Badillo-Urquiola, and Pamela Wisniewski. 2020. Privacy norms within the internet of things using contextual integrity. In *Companion of the 2020 ACM International Conference on Supporting Group Work*. 131–134.
- [4] Jacob Abbott, Jayati Dev, DongInn Kim, Shakthidhar Reddy Gopavaram, Meera Iyer, Shivani Sadam, Shirang Mare, Tatiana Ringenberg, Vafa Andalibi, and L Jean Camp. 2023. Kids, Cats, and Control: Designing Privacy and Security Dashboards for IoT Home Devices. In *Proceedings 2023 Symposium on Usable Security. Internet Society*.
- [5] 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.
- [6] Noura Abdi, Xiao Zhan, Kopo M Ramokapane, and Jose Such. 2021. Privacy norms for smart home personal assistants. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
- [7] Chon Abraham, Dave Chatterjee, and Ronald R. Sims. 2019. Muddling through cybersecurity: Insights from the U.S. healthcare industry. *Business Horizons* 62, 4 (July 2019), 539–548. <https://doi.org/10.1016/j.bushor.2019.03.010>
- [8] Lawrence Abrams. 2022. QNAP force-installs update after DeadBolt ransomware hits 3,600 devices. <https://www.bleepingcomputer.com/news/security/qnap-force-installs-update-after-deadbolt-ransomware-hits-3-600-devices/>
- [9] Yasemin Acar, Sascha Fahl, and Michelle L Mazurek. 2016. You are not your developer, either: A research agenda for usable security and privacy research beyond end users. *2016 IEEE Cybersecurity Development (SecDev)* (2016), 3–8.
- [10] Sara Amini and Chris Kanich. 2017. Characterizing the impact of malware infections and remediation attempts through support forum analysis. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*. 70–78. <https://doi.org/10.1109/ECRIME.2017.7945056>

- [11] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (July 2018), 1–23. <https://doi.org/10.1145/3214262>
- [12] Katrin Auspurg and Thomas Hinz. 2014. *Factorial survey experiments*. Vol. 175. Sage Publications.
- [13] An Baeyens and Tom Goffin. 2015. Boston Scientific Medizintechnik GmbH v. AOK Sachsen-Anhalt. *European journal of health law* 22, 3 (2015), 301–307. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62013CJ0503>
- [14] Paritosh Bahirat, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. 2021. Overlooking Context: How do Defaults and Framing Reduce Deliberation in Smart Home Privacy Decision-Making?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 1–18. <https://doi.org/10.1145/3411764.3445672>
- [15] Natã M Barbosa, Zhuohao Zhang, and Yang Wang. 2020. Do Privacy and Security Matter to Everyone? Quantifying and Clustering User-Centric Considerations About Smart Home Device Adoption. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. 417–435.
- [16] David Barrera, Christopher Bellman, and Paul C van Oorschot. 2022. Security Best Practices: A Critical Analysis Using IoT as a Case Study. *ACM Transactions on Privacy and Security* (2022).
- [17] A. Barth, A. Datta, J.C. Mitchell, and H. Nissenbaum. 2006. Privacy and contextual integrity: framework and applications. In *2006 IEEE Symposium on Security and Privacy (S&P'06)*. 15 pp.–198. <https://doi.org/10.1109/SP.2006.32>
- [18] BBC. 2020. Sonos U-turn over 'bricking' its smart speakers. <https://www.bbc.com/news/technology-51768574>
- [19] Adam Beautement, Ingolf Becker, Simon Parkin, Kat Krol, and Angela Sasse. 2016. Productive security: A scalable methodology for analysing employee security behaviours. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. 253–270.
- [20] Ingolf Becker, Simon Parkin, and M Angela Sasse. 2017. Measuring the Success of {Context-Aware} Security Behaviour Surveys. In *The LASER Workshop: Learning from Authoritative Security Experiment Results (LASER 2017)*. 77–86.
- [21] Behavioural Economics Team of the Australian Government (BETA). 2022. Stay Smart: Helping consumers choose cyber secure smart devices. Australian Government. <https://behaviouraleconomics.pmc.gov.au/sites/default/files/projects/beta-report-cyber-security-labels.pdf>

- [22] Jeff Bennett and Vic Adamowicz. 2001. Some fundamentals of environmental choice modelling. *The choice modelling approach to environmental valuation* (2001), 37–69.
- [23] John M. Blythe, Lynne Coventry, and Linda Little. 2015. Unpacking security policy compliance: the motivators and barriers of employees' security behaviors. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security (SOUPS '15)*. USENIX Association, 103–122.
- [24] 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.
- [25] John M Blythe, Nissy Sombatruang, and Shane D Johnson. 2019. What security features and crime prevention advice is communicated in consumer IoT device manuals and support pages? *Journal of Cybersecurity* 5, 1 (2019), tyz005.
- [26] Tamara Bondar, Hala Assal, and AbdelRahman Abdou. 2023. Why do Internet Devices Remain Vulnerable? A Survey with System Administrators. *NDSS Symposium* (2023). <https://www.ndss-symposium.org/ndss-paper/auto-draft-421/>
- [27] Brennen Bouwmeester, Elsa Rodríguez, Carlos Gañán, Michel van Eeten, and Simon Parkin. 2021. "The Thing Doesn't Have a Name": Learning from Emergent Real-World Interventions in Smart Home Security. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*. 493–512.
- [28] Nellie Bowles. 2018. Thermostats, Locks and Lights: Digital Tools of Domestic Abuse. <https://www.nytimes.com/2018/06/23/technology/smart-home-devices-domestic-abuse.html>
- [29] Anu Bradford. 2020. *The Brussels Effect: How the European Union Rules the World*. Oxford University Press.
- [30] Conner Bradley and David Barrera. 2023. Escaping Vendor Mortality: A New Paradigm for Extending IoT Device Longevity. In *Proceedings of the 2023 New Security Paradigms Workshop (NSPW '23)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3633500.3633501>
- [31] 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.
- [32] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.
- [33] Justin Brookman. 2025. Consumer Reports, US PIRG, Secure Resilient Future Foundation, and the Center for Democracy and Technology Propose, "Connected Consumer Products End of Life Disclosure Act" to Address IoT Security Risks.

https://advocacy.consumerreports.org/press_release/consumer-reports-us-pirg-and-secure-resilient-future-foundation-propose-connected-consumer-products-end-of-life-disclosure-act-to-address-iot-security-risks/

- [34] Michael Brown. 2020. Ring announces 6 new products, along with measures to quiet criticism of its privacy practices. <https://www.techhive.com/article/584254/ring-announces-6-new-products-along-with-measures-to-quiet-criticism-of-its-privacy-practices.html>
- [35] Jonathan Bundy and Michael D. Pfarrer. 2015. A Burden of Responsibility: The Role of Social Approval at the Onset of a Crisis. *Academy of Management Review* 40, 3 (July 2015), 345–369. <https://doi.org/10.5465/amr.2013.0027>
- [36] Jonathan Bundy, Michael D. Pfarrer, Cole E. Short, and W. Timothy Coombs. 2017. Crises and Crisis Management: Integration, Interpretation, and Research Development. *Journal of Management* 43, 6 (July 2017), 1661–1692. <https://doi.org/10.1177/0149206316680030>
- [37] Zach Capers. 2022. More Healthcare Devices Means More Cyberattacks — How Weak Medical IoT Security Threatens Patient Care. <https://www.capterra.com/resources/medical-internet-of-things-iot-security/>
- [38] Seth Carmody, Andrea Coravos, Ginny Fahs, Audra Hatch, Janine Medina, Beau Woods, and Joshua Corman. 2021. Building resilient medical technology supply chains with a software bill of materials. *npj Digital Medicine* 4, 1 (Feb. 2021), 1–6. <https://doi.org/10.1038/s41746-021-00403-w> Number: 1 Publisher: Nature Publishing Group.
- [39] Nico Castelli, Corinna Ogonowski, Timo Jakobi, Martin Stein, Gunnar Stevens, and Volker Wulf. 2017. What happened in my home? An end-user development approach for smart home data visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 853–866.
- [40] Jason Ceci, Jonah Stegman, and Hassan Khan. 2023. No privacy in the electronics repair industry. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3347–3364.
- [41] Microsoft Security Response Center. 2019. Corporate IoT - a path to intrusion. <https://www.microsoft.com/en-us/msrc/blog/2019/08/corporate-iot-a-path-to-intrusion/>
- [42] Orçun Çetin, Carlos Ganán, Lisette Altena, Takahiro Kasama, Daisuke Inoue, Kazuki Tamiya, Ying Tie, Katsunari Yoshioka, and Michel Van Eeten. 2019. Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai.. In *NDSS*.
- [43] George Chalhoub, Ivan Flechais, Norbert Nthala, and Ruba Abu-Salma. 2020. Innovation Inaction or In Action? The Role of User Experience in the Security and Privacy Design of Smart Home Cameras. 185–204.

- [44] 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*. ACM, 1–16. <https://doi.org/10.1145/3411764.3445691>
- [45] G. Chalhoub and A. Martin. 2023. But is it exploitable? Exploring how Router Vendors Manage and Patch Security Vulnerabilities in Consumer-Grade Routers. *European Symposium on Usable Security (EuroUSEC 2023)* (2023), 277–295.
- [46] Ramaswamy Chandramouli. 2022. Guide to a Secure Enterprise Network Landscape. *NIST Special Publication* (2022).
- [47] Jiahong Chen and Lachlan Urquhart. 2022. ‘They’re all about pushing the products and shiny things rather than fundamental security’: Mapping socio-technical challenges in securing the smart home. *Information & Communications Technology Law* 31, 1 (2022), 99–122.
- [48] Catalin Cimpanu. 2020. Hackers are hijacking smart building access systems to launch DDoS attacks. <https://www.zdnet.com/article/hackers-are-hijacking-smart-building-access-systems-to-launch-ddos-attacks/>
- [49] European Commission. 2025. C(2025)618 – Standardisation request M/606. https://ec.europa.eu/growth/tools-databases/enorm/mandate/606_en
- [50] European Commission. 2025. EU cybersecurity certification framework. <https://digital-strategy.ec.europa.eu/en/policies/cybersecurity-certification-framework>
- [51] European Commission. 2025. EU Cybersecurity Strategy. <https://digital-strategy.ec.europa.eu/en/policies/cybersecurity-strategy>
- [52] Lynne Coventry and Dawn Branley. 2018. Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. *Maturitas* 113 (July 2018), 48–52. <https://doi.org/10.1016/j.maturitas.2018.04.008>
- [53] Lynne Coventry, Dawn Branley-Bell, Elizabeth Sillence, Sabina Magalini, Pasquale Mari, Aimilia Magkanaraki, and Kalliopi Anastasopoulou. 2020. Cyber-Risk in Healthcare: Exploring Facilitators and Barriers to Secure Behaviour. In *International conference on human-computer interaction*. Springer International Publishing, 105–122. https://doi.org/10.1007/978-3-030-50309-3_8
- [54] Jayne Cox, Sarah Griffith, Sara Giorgi, and Geoff King. 2013. Consumer understanding of product lifetimes. *Resources, Conservation and Recycling* 79 (2013), 21–29.
- [55] Ry Crist. 2021. ADT technician pleads guilty to spying on customer camera feeds for years. <https://www.cnet.com/home/smart-home/adt-home-security-technician-pleads-guilty-to-spying-on-customer-camera-feeds-for-years/>

- [56] Cynerio. 2022. Cynerio Research Finds Critical Medical Device Risks Continue to Threaten Hospital Security and Patient Safety. <https://www.cynerio.com/blog/cynerio-research-finds-critical-medical-device-risks-continue-to-threaten-hospital-security-and-patient-safety>
- [57] 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.
- [58] 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.
- [59] Wendy Davis. 2015. Vizio Hit With Privacy Lawsuit Over Connected TVs. <https://www.mediapost.com/publications/article/262835/vizio-hit-with-privacy-lawsuit-over-connected-tvs.html>
- [60] Mattt Day, Giles Turner, and Natalia Drozdiak. 2019. Thousands of Amazon Workers Listen to Alexa Users' Conversations. <https://time.com/5568815/amazon-workers-listen-to-alexa/>
- [61] Benjamin Dean. 2018. *An Exploration of Strict Products Liability and the Internet of Things*. Technical Report. Center for Democracy and Technology. <https://cdt.org/wp-content/uploads/2018/04/2018-04-16-IoT-Strict-Products-Liability-FNL.pdf>
- [62] DeepL. 2024. DeepL translator. <https://www.deepl.com/en/translator>
- [63] Department for Digital, Culture, Media and Sport (UK Government). 2018. Code of Practice for Consumer IoT Security. (2018). <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security>
- [64] Nesara Dissanayake, Asangi Jayatilaka, Mansooreh Zahedi, and M. Ali Babar. 2022. Software security patch management - A systematic literature review of challenges, approaches, tools and practices. *Information and Software Technology* 144 (April 2022), 106771. <https://doi.org/10.1016/j.infsof.2021.106771>
- [65] Nesara Dissanayake, Asangi Jayatilaka, Mansooreh Zahedi, and Muhammad Ali Babar. 2023. An Empirical Study of Automation in Software Security Patch Management. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*. Association for Computing Machinery, 1–13. <https://doi.org/10.1145/3551349.3556969>
- [66] Nesara Dissanayake, Mansooreh Zahedi, Asangi Jayatilaka, and Muhammad Ali Babar. 2021. A grounded theory of the role of coordination in software security patch management. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2021)*. Association for Computing Machinery, 793–805. <https://doi.org/10.1145/3468264.3468595>

- [67] Nesara Dissanayake, Mansooreh Zahedi, Asangi Jayatilaka, and Muhammad Ali Babar. 2022. Why, How and Where of Delays in Software Security Patch Management: An Empirical Investigation in the Healthcare Sector. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 362:1–362:29. <https://doi.org/10.1145/3555087>
- [68] Paul Dourish, Rebecca E Grinter, Jessica Delgado De La Flor, and Melissa Joseph. 2004. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing* 8 (2004), 391–401.
- [69] Guillaume Dupont, Daniel Ricardo dos Santos, Elisa Costante, Jerry den Hartog, and Sandro Etalle. 2020. A Matter of Life and Death: Analyzing the Security of Healthcare Networks. In *ICT Systems Security and Privacy Protection (IFIP Advances in Information and Communication Technology)*. Springer International Publishing, Cham, 355–369. https://doi.org/10.1007/978-3-030-58201-2_24
- [70] 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)*. 447–464. <https://doi.org/10.1109/SP40000.2020.00043>
- [71] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2021. Which Privacy and Security Attributes Most Impact Consumers’ Risk Perception and Willingness to Purchase IoT Devices?. In *2021 IEEE Symposium on Security and Privacy (SP)*. 519–536. <https://doi.org/10.1109/SP40001.2021.00112>
- [72] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of {IoT} Devices? 1505–1522. <https://www.usenix.org/conference/usenixsecurity23/presentation/emami-naeini>
- [73] 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*. ACM, 1–12. <https://doi.org/10.1145/3290605.3300764>
- [74] ENISA. 2019. Good Practices for Security of IoT - Secure Software Development Lifecycle. (2019). <https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot-1>
- [75] Michael Fagan, Mohammad Maifi Hasan Khan, and Ross Buck. 2015. A study of users’ experiences and beliefs about software update messages. *Computers in Human Behavior* 51 (Oct. 2015), 504–519. <https://doi.org/10.1016/j.chb.2015.04.075>
- [76] Michael Fagan, Katerina Megias, Karen Scarfone, and Matthew Smith. 2020. *Foundational Cybersecurity Activities for IoT Device Manufacturers*. Technical Report

- NIST Internal or Interagency Report (NISTIR) 8259. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.IR.8259>
- [77] Sadeqh Farhang, Jake Weidman, Mohammad Mahdi Kamani, Jens Grossklags, and Peng Liu. 2018. Take It or Leave It: A Survey Study on Operating System Upgrade Practices. In *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC '18)*. Association for Computing Machinery, New York, NY, USA, 490–504. <https://doi.org/10.1145/3274694.3274733>
- [78] Jean Faugier and Mary Sargeant. 1997. Sampling hard to reach populations. *Journal of advanced nursing* 26, 4 (1997), 790–797.
- [79] Nathan Favero and Minjung Kim. 2021. Everything Is Relative: How Citizens Form and Use Expectations in Evaluating Services. *Journal of Public Administration Research and Theory* 31, 3 (July 2021), 561–577. <https://doi.org/10.1093/jopart/muaa048>
- [80] Federal Communications Commission (FCC). 2024. U.S. Cyber Trust Mark. <https://www.fcc.gov/CyberTrustMark#:~:text=The%20program%20applies%20to%20consumer,door%20openers%2C%20and%20baby%20monitors.>
- [81] Federal Trade Commission (FTC). 2024. Smart Device Makers' Failure to Provide Updates May Leave You Smarting. <https://www.ftc.gov/reports/smart-device-makers-failure-provide-updates-may-leave-you-smarting>
- [82] United Nations Economic Commission for Europe (UNECE). 2021. UN Regulation No. 155 - Cyber security and cyber security management system. <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security>
- [83] Federal Office for Information Security (BSI). 2025. The Label for IT Security in the Consumer Market. https://www.bsi.bund.de/EN/Themen/Verbraucherinnen-und-Verbraucher/IT-SiK-fuer-Verbraucher/IT-SiK-fuer-Verbraucher_node.html
- [84] District Court for the Western District of Pennsylvania. 2022. Cohen v. Johnson & Johnson. <https://law.justia.com/cases/federal/district-courts/pennsylvania/pawdce/2:2020cv00057/262978/68/>
- [85] International Medical Device Regulators Forum. 2020. Principles and Practices for Medical Device Cybersecurity. <https://www.imdrf.org/documents/principles-and-practices-medical-device-cybersecurity>
- [86] International Medical Device Regulators Forum. 2023. Principles and Practices for the Cybersecurity of Legacy Medical Devices. <https://www.imdrf.org/documents/principles-and-practices-cybersecurity-legacy-medical-devices>

- [87] Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman. 2019. Privacy and security threat models and mitigation strategies of older adults. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 21–40.
- [88] Sandra Gabriele and Sonia Chiasson. 2020. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 1–12. <https://doi.org/10.1145/3313831.3376651>
- [89] Vaibhav Garg. 2021. A Lemon by Any Other Label. *ICISSP (2021)*, 558–565.
- [90] Vaibhav Garg and Andreas Kuehn. 2021. Squeezing the Cybersecurity Lemons – A Labeling Regime for IoT Products. *Usenix ;login:* (2021). <https://www.usenix.org/publications/loginonline/squeezing-cybersecurity-lemons-%E2%80%93-labeling-regime-iot-products>
- [91] Simon Geiregat. 2022. What Digital Content Consumers (Should) Want. In *Immaterialgüter und Medien im Binnenmarkt*, Anna K. Bernzen, Karina Grisse, and Katharina Kaesling (Eds.). Nomos Verlagsgesellschaft mbH & Co. KG, 65–88. <https://doi.org/10.5771/9783748934233-65>
- [92] Alex Gnanapragasam, Christine Cole, Jagdeep Singh, and Tim Cooper. 2018. Consumer perspectives on longevity and reliability: a national study of purchasing factors across eighteen product categories. *Procedia Cirp* 69 (2018), 910–915.
- [93] Alex Gnanapragasam, Masahiro Oguchi, Christine Cole, and Tim Cooper. 2017. Consumer expectations of product lifetimes around the world: a review of global research findings and methods. *PLATE: Product Lifetimes And The Environment* (2017), 464–469.
- [94] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [95] Shakthidhar Gopavaram, Jayati Dev, Sanchari Das, and L Jean 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)*.
- [96] Government of the United Kingdom. 2002. The Medical Devices Regulations 2002 (SI 2002/618). <https://www.legislation.gov.uk/uksi/2002/618/contents/made>
- [97] Andy Greenberg. 2015. After Jeep Hack, Chrysler Recalls 1.4M Vehicles for Bug Fix. *Wired* (2015). <https://www.wired.com/2015/07/jeep-hack-chrysler-recalls-1-4m-vehicles-bug-fix/>
- [98] Ames Gross. 2023. Japan Outlines New Medical Device Cybersecurity Regulation. <https://www.pacificbridgemedical.com/news-brief/japan-outlines-new-medical-device-cybersecurity-regulation/>

- [99] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How Many Interviews Are Enough?: An Experiment with Data Saturation and Variability. *Field methods* 18, 1 (2006), 59–82. <https://doi.org/10.1177/1525822X05279903>
- [100] Marco Gutfleisch, Markus Schöps, Jonas Hielscher, Mary Cheney, Sibel Sayin, Nathalie Schuhmacher, Ali Mohamad, and M. Angela Sasse. 2022. Caring About IoT-Security – An Interview Study in the Healthcare Sector. In *Proceedings of the 2022 European Symposium on Usable Security (EuroUSEC '22)*. Association for Computing Machinery, 202–215. <https://doi.org/10.1145/3549015.3554209>
- [101] Julie Haney, Yasemin Acar, and Susanne Furman. 2021. "It's the Company, the Government, You and I": User Perceptions of Responsibility for Smart Home Privacy and Security. In *30th USENIX Security Symposium (USENIX Security 21)*. 411–428.
- [102] 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. https://doi.org/10.1007/978-3-030-79318-0_6
- [103] Julie M. Haney and Susanne M. Furman. 2023. Smart Home Device Loss of Support: Consumer Perspectives and Preferences. In *HCI for Cybersecurity, Privacy and Trust: 5th International Conference, HCI-CPT 2023, Held as Part of the 25th HCI International Conference, HCII 2023, Copenhagen, Denmark, July 23–28, 2023, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 492–510. https://doi.org/10.1007/978-3-031-35822-7_32
- [104] Julie M. Haney and Susanne M. Furman. 2023. User Perceptions and Experiences with Smart Home Updates. In *2023 IEEE Symposium on Security and Privacy (SP)*. 2867–2884. <https://doi.org/10.1109/SP46215.2023.10179459>
- [105] Julie M. Haney, Susanne M. Furman, and Yasemin Acar. 2020. Smart Home Security and Privacy Mitigations: Consumer Perceptions, Practices, and Challenges. In *HCI for Cybersecurity, Privacy and Trust (Lecture Notes in Computer Science)*, Abbas Moallem (Ed.). Springer International Publishing, 393–411. https://doi.org/10.1007/978-3-030-50309-3_26
- [106] Scharon Harding. 2025. Bose SoundTouch home theater systems regress into dumb speakers Feb. 18. <https://arstechnica.com/gadgets/2025/10/bose-soundtouch-home-theater-systems-regress-into-dumb-speakers-feb-18/>
- [107] Scharon Harding. 2025. Logitech will brick its \$ 100 Pop smart home buttons on October 15. <https://arstechnica.com/gadgets/2025/10/logitech-will-brick-its-100-pop-smart-home-buttons-on-october-15/>
- [108] Scharon Harding. 2025. People regret buying Amazon smart displays after being bombarded with ads. <https://arstechnica.com/gadgets/2025/10/people->

regret-buying-amazon-smart-displays-after-being-bombarded-with-ads/

- [109] Harris Interactive. 2019. Consumer Internet of Things Security Labelling Survey Research Findings. Consultation on regulatory proposals on consumer IoT security, UK Department of Culture, Media, and Sport (DCMS). https://assets.publishing.service.gov.uk/media/5ff713bfe90e0763a31280a1/Harris_Interactive_Consumer_IoT_Security_Labelling_Survey_Report_V2.pdf
- [110] H. L. A. Hart. 2012. *The concept of law* (3 ed.). Oxford University Press, London, England.
- [111] Ayako A Hasegawa, Naomi Yamashita, Tatsuya Mori, Daisuke Inoue, and Mitsuaki Akiyama. 2022. Understanding Non-Experts' Security-and Privacy-Related Questions on a Q&A Site. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 39–56.
- [112] Center for Devices and Radiological Health. 2016. Postmarket Management of Cybersecurity in Medical Devices. *U.S. Food and Drug Administration* (Dec. 2016). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/postmarket-management-cybersecurity-medical-devices>
- [113] Center for Devices and Radiological Health. 2023. Cybersecurity in Medical Devices: Refuse to Accept Policy for Cyber Devices and Related Systems Under Section 524B of the FD&C Act. *U.S. Food and Drug Administration* (March 2023). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-refuse-accept-policy-cyber-devices-and-related-systems-under-section>
- [114] Center for Devices and Radiological Health and Center for Biologics Evaluation and Research. 2022. Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions. *U.S. Food and Drug Administration* (April 2022). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions>
- [115] Securin Health-ISAC, Finite State. 2023. Exploitable Vulnerabilities That Expose Healthcare Facilities Surged Nearly 60% Since 2022, New Research Report Finds. *Health-ISAC - Health Information Sharing and Analysis Center* (Aug. 2023). <https://h-isac.org/2023-state-of-cybersecurity-for-medical-devices-and-healthcare-systems/>
- [116] James J Heckman. 1990. Selection Bias and Self-selection. In *Econometrics*. Springer, 201–224.
- [117] Bronwyn Higgs, Michael Jay Polonsky, and Mary Hollick. 2005. Measuring expectations: forecast vs. ideal expectations. Does it really matter? *Journal of retailing and consumer services* 12, 1 (2005), 49–64.

- [118] Microsoft Threat Intelligence. 2023. IoT devices and Linux-based systems targeted by OpenSSH trojan campaign. <https://www.microsoft.com/en-us/security/blog/2023/06/22/iot-devices-and-linux-based-systems-targeted-by-openssh-trojan-campaign/>
- [119] International Electrotechnical Commission (IEC). 2021. IEC 81001-5-1:2021. <https://www.iso.org/standard/76097.html>
- [120] Mohammad S. Jalali and Jessica P. Kaiser. 2018. Cybersecurity in Hospitals: A Systematic, Organizational Perspective. *Journal of Medical Internet Research* 20, 5 (2018), e10059. <https://doi.org/10.2196/10059>
- [121] Adam Jenkins, Pieris Kalligeros, Kami Vaniea, and Maria K Wolters. 2020. “Anyone Else Seeing this Error?”: Community, System Administrators, and Patch Information. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 105–119.
- [122] Sijia Jiang and Jim Finkle. 2016. China’s Xionghai to recall up to 10,000 webcams after hack. *Reuters* (Oct. 2016). <https://www.reuters.com/article/us-cyber-attacks-china-idUSKCN12P1TT>
- [123] Haojian Jin, Boyuan Guo, Rituparna Roychoudhury, Yaxing Yao, Swarun Kumar, Yuvraj Agarwal, and Jason I. Hong. 2022. Exploring the Needs of Users for Supporting Privacy-Protective Behaviors in Smart Homes. In *CHI Conference on Human Factors in Computing Systems*. ACM, 1–19. <https://doi.org/10.1145/3491102.3517602>
- [124] 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.
- [125] Laura Kocksch, Matthias Korn, Andreas Poller, and Susann Wagenknecht. 2018. Caring for IT security: Accountabilities, moralities, and oscillations in IT security practices. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–20.
- [126] Constantinos Koliadis, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. 2017. DDoS in the IoT: Mirai and other botnets. *Computer* 50, 7 (2017), 80–84.
- [127] Ross Koppel, Sean Smith, Jim Blythe, and Vijay Kothari. 2015. Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient? In *Driving Quality in Informatics: Fulfilling the Promise*. IOS Press, 215–220. <https://doi.org/10.3233/978-1-61499-488-6-215>
- [128] Anastassija Kostan, Sara Olschar, Lucy Simko, and Yasemin Acar. 2024. Exploring digital security and privacy in relative poverty in Germany through qualitative interviews. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2029–2046.

- [129] Brian Krebs. 2025. DDoS Botnet Aisuru Blankets US ISPs in Record DDoS. <https://krebsonsecurity.com/2025/10/ddos-botnet-aisuru-blankets-us-isps-in-record-ddos/>
- [130] Kat Krol, Jonathan M Spring, Simon Parkin, and M Angela Sasse. 2016. Towards robust experimental design for user studies in security and privacy. In *The LASER Workshop: Learning from Authoritative Security Experiment Results (LASER 2016)*. 21–31.
- [131] 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.
- [132] Lorenz Kustosch, Carlos Gañán, Michel van Eeten, and Simon Parkin. 2025. Patching Up: Stakeholder Experiences of Security Updates for Connected Medical Devices. In *34th USENIX Security Symposium (USENIX Security 25)*. USENIX Association, 2265–2281. <https://www.usenix.org/conference/usenixsecurity25/presentation/kustosch-patching>
- [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)*. USENIX Association, Anaheim, CA, 1487–1504. <https://www.usenix.org/conference/usenixsecurity23/presentation/kustosch>
- [134] Lorenz Kustosch, Carlos Gañán, Mattis van't Schip, Michel van Eeten, and Simon Parkin. 2025. Regulating Smart Device Support Periods: User Expectations and the European Cyber Resilience Act. In *34th USENIX Security Symposium (USENIX Security 25)*. USENIX Association, 5149–5168. <https://www.usenix.org/conference/usenixsecurity25/presentation/kustosch-regulating>
- [135] Forescout Research Labs. 2020. Connected Medical Device Security: A Deep Dive into Healthcare Networks. <https://www.forescout.com/resources/connected-medical-device-security-a-deep-dive-into-healthcare-networks/>
- [136] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction 2*, CSCW (Nov. 2018), 1–31. <https://doi.org/10.1145/3274371>
- [137] Ying-Tsung Lee, Wei-Hsuan Hsiao, Yan-Shao Lin, and Seng-Cho T Chou. 2017. Privacy-preserving data analytics in cloud-based smart home with community hierarchy. *IEEE Transactions on Consumer Electronics* 63, 2 (2017), 200–207.

- [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] Ye Li, Antonia Krefeld-Schwalb, Daniel G. Wall, Eric J. Johnson, Olivier Toubia, and Daniel M. Bartels. 2021. The More You Ask, the Less You Get: When Additional Questions Hurt External Validity. *Journal of Marketing Research* (Dec. 2021). <https://doi.org/10.1177/002224372111073581>
- [140] Peiyu Liu, Shouling Ji, Lirong Fu, Kangjie Lu, Xuhong Zhang, Jingchang Qin, Wenhai Wang, and Wenzhi Chen. 2023. How IoT Re-using Threatens Your Sensitive Data: Exploring the User-Data Disposal in Used IoT Devices. In *2023 IEEE Symposium on Security and Privacy (SP)*. 3365–3381. <https://doi.org/10.1109/SP46215.2023.10179294> ISSN: 2375-1207.
- [141] Kaspar Rosager Ludvigsen. 2023. The Role of Cybersecurity in Medical Devices Regulation: Future Considerations and Solutions. *Law, Technology and Humans* 5, 2 (Nov. 2023), 59–77. <https://doi.org/10.5204/1thj.3080>
- [142] René Mahieu, Hadi Asghari, Christopher Parsons, Joris van Hoboken, Masashi Crete-Nishihata, Andrew Hiltz, and Siena Anstis. 2021. Measuring the Brussels Effect through Access Requests: Has the European General Data Protection Regulation Influenced the Data Protection Rights of Canadian Citizens? *Journal of Information Policy* 11 (Dec. 2021), 301–349. <https://doi.org/10.5325/jinfopoli.11.2021.0301>
- [143] Prianka Mandal, Amit Seal Ami, Iria Giuffrida, Daniel Shin, Ella Sullivan, and Adwait Nadkarni. 2025. “We can’t Allow IoT Vendors to Pass off all Such Liability to the Consumer”: Investigating the U.S. Legal Perspectives on Liability for IoT Product Security. In *2025 IEEE Symposium on Security and Privacy (SP)*. 3746–3764. <https://doi.org/10.1109/SP61157.2025.00212>
- [144] Prianka Mandal, Amit Seal Ami, Victor Olaiya, Sayyed Hadi Razmjo, and Adwait Nadkarni. 2024. “Belt and suspenders” or “just red tape”? Investigating Early Artifacts and User Perceptions of {IoT} App Security Certification. 4927–4944. <https://www.usenix.org/conference/usenixsecurity24/presentation/mandal>
- [145] Prianka Mandal and Adwait Nadkarni. 2024. “We can’t change it overnight”: Understanding Industry Perspectives on IoT Product Security Compliance and Certification. *IEEE Computer Society*, 91–91. <https://doi.org/10.1109/SP61157.2025.00091> ISSN: 2375-1207.
- [146] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. 2020. “I don’t know how to protect myself”: Understanding Privacy Perceptions Resulting from the Presence of Bystanders in Smart Environments. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–11.

- [147] Kirsten Martin. 2015. Privacy notices as tabula rasa: An empirical investigation into how complying with a privacy notice is related to meeting privacy expectations online. *Journal of Public Policy & Marketing* 34, 2 (2015), 210–227.
- [148] Kirsten Martin and Katie Shilton. 2016. Why experience matters to privacy: How context-based experience moderates consumer privacy expectations for mobile applications. *Journal of the Association for Information Science and Technology* 67, 8 (2016), 1871–1882. <https://doi.org/10.1002/asi.23500>
- [149] Clayton J Masterman and W Kip Viscusi. 2020. The Specific Consumer Expectations Test for Product Defects. *Ind. LJ* 95 (2020), 183.
- [150] Arunesh Mathur, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. 2018. Quantifying Users’ Beliefs about Software Updates. In *NDSS Workshop on Usable Security*. <https://doi.org/10.14722/usec.2018.23036>
- [151] Emma McMahon, Ryan Williams, Malaka El, Sagar Samtani, Mark Patton, and Hsinchun Chen. 2017. Assessing medical device vulnerabilities on the Internet of Things. In *2017 IEEE International Conference on Intelligence and Security Informatics (ISI)*. 176–178. <https://doi.org/10.1109/ISI.2017.8004903>
- [152] Megan Crouse. 2025. WUS to Launch Cyber Trust Mark to Label Secure Smart Devices. <https://www.techrepublic.com/article/us-cyber-trust-mark-iot-security/>
- [153] Joseph Menn, Jim Finkle, and Dustin Volz. 2016. Cyber attacks disrupt PayPal, Twitter, other sites. *Reuters* (Oct. 2016). <https://www.reuters.com/article/us-usa-cyber-idUSKCN12L1ME>
- [154] Sharan B. Merriam and Elizabeth J. Tisdell. 2015. *Qualitative Research: A Guide to Design and Implementation*. John Wiley & Sons. Google-Books-ID: JFN_BwAAQBAJ.
- [155] Louis Meuleman. 2008. *Public management and the metagovernance of hierarchies, networks and markets: The feasibility of designing and managing governance style combinations*. Springer. https://doi.org/10.1007/978-3-7908-2054-6_4
- [156] MHRA. 2024. Implementation of the Future Regulations. <https://www.gov.uk/government/publications/implementation-of-the-future-regulation-of-medical-devices/implementation-of-the-future-regulations>
- [157] MHRA. 2025. Alerts, recalls and safety information: drugs and medical devices. https://www.gov.uk/drug-device-alerts?alert_type%5B%5D=field-safety-notices
- [158] Trade Ministry of Economy and Industry. 2024. IoT Product Security Conformity Assessment Scheme Policy Draft. https://www.meti.go.jp/english/press/2024/0315_001.html

- [159] Yisroel Mirsky, Tom Mahler, Ilan Shelef, and Yuval Elovici. 2019. CT-GAN: Malicious Tampering of 3D Medical Imagery using Deep Learning.. In *USENIX Security Symposium*, Vol. 2019.
- [160] 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 (SP)*. 429–446. <https://doi.org/10.1109/SP40000.2020.00021>
- [161] Mozilla. 2021. *Privacy not included. <https://foundation.mozilla.org/en/privacynotincluded/>. Accessed: 2021-05-25.
- [162] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujó Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy expectations and preferences in an {IoT} world. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. 399–412.
- [163] Yuhong Nan, Xueqiang Wang, Luyi Xing, Xiaojing Liao, Ruoyu Wu, Jianliang Wu, Yifan Zhang, and XiaoFeng Wang. 2023. Are you spying on me? Large-Scale analysis on IoT data exposure through companion apps. In *32nd USENIX Security Symposium (USENIX Security 23)*. 6665–6682.
- [164] 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.
- [165] RVO Netherlands Enterprise Agency. 2022. Supplying updates is mandatory. (2022). <https://business.gov.nl/amendment/supplying-updates-to-become-mandatory/>
- [166] BBC News. 2019. Google seeks permission for staff to listen to Assistant recordings. *BBC News* (Sept. 2019). <https://www.bbc.com/news/technology-49796207>
- [167] NHS. 2023. Cyber security guidance for healthcare professionals procuring and deploying connected medical devices. <https://digital.nhs.uk/cyber-and-data-security/guidance-and-assurance/guidance-for-procuring-and-deploying-connected-medical-devices>
- [168] NHS. 2025. Cyber alerts. <https://digital.nhs.uk/cyber-alerts>
- [169] Laura Lynggaard Nielsen. 2022. What Makes IoT Secure? A Maturity Analysis of Industrial Product Manufacturers’ Approaches to IoT Security. In *International Conference on Human-Computer Interaction*. Springer, 406–421.
- [170] Daisuke Nishijima and Masahiro Oguchi. 2023. Measuring product lifetime extension potential by increasing the expected product lifetime: Methodology and case study. *Business Strategy and the Environment* 32, 4 (2023), 1218–1231.
- [171] NIST. 2025. internet of things. https://csrc.nist.gov/glossary/term/internet_of_things

- [172] Geoff Norman. 2010. Likert scales, levels of measurement and the “laws” of statistics. *Advances in Health Sciences Education* 15, 5 (Dec. 2010), 625–632. <https://doi.org/10.1007/s10459-010-9222-y>
- [173] 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.
- [174] Dawn Carla Nunziato. 2023. The Digital Services Act and the Brussels Effect on Platform Content Moderation. *Chi. J. Int'l L.* 24 (2023), 115. <https://cjil.uchicago.edu/print-archive/digital-services-act-and-brussels-effect-platform-content-moderation>
- [175] Lindsey O'Donnell. 2019. EU Recalls Children's Smartwatch That Leaks Location Data. <https://threatpost.com/eu-recalls-childrens-smartwatch-that-leaks-location-data/141511/>
- [176] Lindsey O'Donnell. 2020. Unpatched Security Flaws Open Connected Vacuum to Takeover. <https://threatpost.com/unpatched-security-flaws-open-connected-vacuum-to-takeover/153142/>
- [177] District Court of Appeal of the State of Florida. 2022. Grieco v. Daiho Sangyo, Inc., 11 pages. <https://law.justia.com/cases/florida/fourth-district-court-of-appeal/2022/20-2294.html>
- [178] State of California. 2018. California Consumer Privacy Act. (2018). https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3&part=4.&lawCode=CIV&title=1.81.5
- [179] Federal Bureau of Investigation (FBI). 2022. Unpatched and Outdated Medical Devices Provide Cyber Attack Opportunities. <https://www.aha.org/system/files/media/file/2022/09/fbi-pin-tlp-white-unpatched-and-outdated-medical-devices-provide-cyber-attack-opportunities-sept-12-2022.pdf>
- [180] Cyber Security Agency of Singapore. 2025. About Cybersecurity Labelling Scheme for IoT - CLS(IoT). <https://www.csa.gov.sg/our-programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme/about>
- [181] Masahiro Oguchi, Tomohiro Tasaki, Ichiro Daigo, Tim Cooper, Christine Cole, and Alex Gnanapragasam. 2016. Consumers' expectations for product lifetimes of consumer durables. In *2016 Electronics Goes Green 2016+(EGG)*. IEEE, 1–6.
- [182] Charlie Osborne. 2021. Remote code execution flaw allowed hijack of Motorola Halo+ baby monitors. <https://portswigger.net/daily-swig/remote-code-execution-flaw-allowed-hijack-of-motorola-halo-baby-monitors>

- [183] Charlie Osborne. 2022. Researcher discloses alleged zero-day vulnerabilities in NUUO NVRmini2 recording device. <https://portswigger.net/daily-swig/researcher-discloses-alleged-zero-day-vulnerabilities-in-nuuo-nvrmini2-recording-device>
- [184] Pierluigi Paganini. 2016. Watch out, FLocker Ransomware targets Android smart TVs. <https://securityaffairs.co/wordpress/48383/iot/flocker-ransomware-smart-tvs.html>
- [185] 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.
- [186] Morgen E. Peck. 2011. Medical Devices Are Vulnerable to Hacks, But Risk Is Low Overall - IEEE Spectrum. <https://spectrum.ieee.org/medical-devices-are-vulnerable-to-hacks-but-risk-is-low-overall>
- [187] 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.
- [188] Lutz Preuss, Ralf Barkemeyer, Olivier Gergaud, and Christophe Faugère. 2018. Corporate Scandals: Firm Response Strategies and Subsequent Media Coverage. *Academy of Management Proceedings* 2018, 1 (Aug. 2018), 18369. <https://doi.org/10.5465/AMBPP.2018.18369abstract>
- [189] 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>
- [190] Prolific. 2022. <https://www.prolific.co/>
- [191] Qualtrics. 2023. <https://www.qualtrics.com>
- [192] Qualtrics. 2024. Qualtrics translate survey. <https://www.qualtrics.com/support/survey-platform/survey-module/survey-tools/translate-survey/>
- [193] Emilee Rader and Rick Wash. 2015. Identifying patterns in informal sources of security information. *Journal of Cybersecurity* 1, 1 (2015), 121–144.
- [194] 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*. 1–17.

- [195] Mathias Reimann. 2003. Product Liability in a Global Context: the Hollow Victory of the European Model. *European Review of Private Law* 11, Issue 2 (April 2003), 128–154. <https://doi.org/10.54648/ERPL2003011>
- [196] Dan Robinson. 2023. Nice smart device – how long does it get software updates? https://www.theregister.com/2023/01/16/smart_device_software_support/
- [197] Elsa Rodríguez, Arman Noroozian, Michel van Eeten, and Carlos Gañán. 2021. Super-Spreaders: Quantifying the Role of IoT Manufacturers in Device Infections. In *20th Annual Workshop on the Economics of Information Security (WEIS 2021)*.
- [198] Elsa Rodríguez, Arman Noroozian, Michel van Eeten, and Carlos Gañán. 2021. Superspreaders: Quantifying the role of IoT manufacturers in device infections. In *Workshop on the Economics of Information Security (WEIS)*.
- [199] Elsa Rodríguez, Max Fukkink, Simon Parkin, Michel van Eeten, and Carlos Gañán. 2022. Difficult for Thee, But Not for Me: Measuring the Difficulty and User Experience of Remediating Persistent IoT Malware. In *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*. 392–409. <https://doi.org/10.1109/EuroSP53844.2022.00032>
- [200] Sean Ross. 2025. How Is a Market Failure Corrected? <https://www.investopedia.com/ask/answers/042115/how-market-failure-corrected.asp#:~:text=Market%20failure%20can%20be%20caused,%2C%20subsidies%2C%20and%20trade%20restrictions.>
- [201] Asreen Rostami, Minna Vigren, Shahid Raza, and Barry Brown. 2022. Being Hacked: Understanding Victims’ Experiences of IoT Hacking. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 613–631.
- [202] Takayuki Sasaki, Tomoya Inazawa, Youhei Yamaguchi, Simon Parkin, Michel van Eeten, Katsunari Yoshioka, and Tsutomu Matsumoto. 2025. Am I Infected? Lessons from Operating a Large-Scale IoT Security Diagnostic Service. 1493–1510. <https://www.usenix.org/conference/usenixsecurity25/presentation/sasaki>
- [203] Carsten Sauer, Katrin Auspurg, and Thomas Hinz. 2020. Designing Multi-Factorial Survey Experiments: Effects of Presentation Style (Text or Table), Answering Scales, and Vignette Order. *Methods, data, analyses: a journal for quantitative methods and survey methodology (mda)* 14, 2 (2020), 195–214.
- [204] Bruce Schneier. 2016. We Need to Save the Internet from the Internet of Things. https://www.schneier.com/essays/archives/2016/10/we_need_to_save_the_.html
- [205] Bruce Schneier. 2018. *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company.

- [206] Bruce Schneier. 2022. Ring Gives Videos to Police without a Warrant or User Consent - Schneier on Security. <https://www.schneier.com/blog/archives/2022/08/ring-gives-videos-to-police-without-a-warrant-or-user-consent.html>
- [207] Bruce Schneier. 2025. TP-Link Router Botnet. <https://www.schneier.com/blog/archives/2025/03/tp-link-router-botnet.html>
- [208] Norbert Schwarz, Hans-J Hippler, Brigitte Deutsch, and Fritz Strack. 1985. Response scales: Effects of category range on reported behavior and comparative judgments. *Public opinion quarterly* 49, 3 (1985), 388–395.
- [209] The New York State Senate. 2025. Senate Bill S8507. <https://www.nysenate.gov/legislation/bills/2025/S8507>
- [210] Irina Shklovski, Scott D. Mainwaring, Halla Hrunn Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: perceptions of privacy and mobile app use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2347–2356. <https://doi.org/10.1145/2556288.2557421>
- [211] Yan Shvartzshnaider, Madelyn Rose Sanfilippo, and Noah Apthorpe. 2021. Contextual Integrity as a Gauge for Governing Knowledge Commons. In *Governing Privacy in Knowledge Commons*, Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg (Eds.). Cambridge University Press, 220–244. <https://doi.org/10.1017/9781108749978.010>
- [212] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. 2016. Learning Privacy Expectations by Crowdsourcing Contextual Informational Norms. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing 4* (Sept. 2016), 209–218.
- [213] Aamir Siddiqui. 2025. Samsung confirms its \$1,800+ fridges will start showing you ads. <https://www.androidauthority.com/samsung-confirms-smart-refrigerator-ads-are-coming-3598848/>
- [214] Satyajit Sinha. 2024. State of IoT 2024: Number of connected IoT devices growing 13% to 18.8 billion globally. <https://iot-analytics.com/number-connected-iot-devices/>
- [215] Yuba R. Siwakoti, Manish Bhurtel, Danda B. Rawat, Adam Oest, and RC Johnson. 2024. Your IP Camera Can Be Abused for Payments: A Study of IoT Exploitation for Financial Services Leveraging Shodan and Criminal Infrastructures. *IEEE Transactions on Consumer Electronics* (2024), 1–1. <https://doi.org/10.1109/TCE.2024.3482708> Conference Name: IEEE Transactions on Consumer Electronics.

- [216] Statista. 2025. Internet of Things - Worldwide. <https://www.statista.com/outlook/tmo/internet-of-things/worldwide>
- [217] Elizabeth Stobert, David Barrera, Valérie Homier, and Daniel Kollek. 2020. Understanding Cybersecurity Practices in Emergency Departments. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3313831.3376881>
- [218] Dan Su and Peter M. Steiner. 2020. An Evaluation of Experimental Designs for Constructing Vignette Sets in Factorial Surveys. *Sociological Methods & Research* 49, 2 (May 2020), 455–497. <https://doi.org/10.1177/0049124117746427>
- [219] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. 2019. “I don’t own the data”: End User Perceptions of Smart Home Device Data Practices and Risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 435–450.
- [220] Madiha Tabassum, Tomasz Kosiński, Alisa Frik, Nathan Malkin, Primal Wijesekera, Serge Egelman, and Heather Richter Lipford. 2019. Investigating Users’ Preferences and Expectations for Always-Listening Voice Assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (Dec. 2019), 1–23. <https://doi.org/10.1145/3369807>
- [221] Mohammad Tahaei and Kami Vaniea. 2019. A survey on developer-centred security. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 129–138.
- [222] 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* 5, 3 (2021), 431–450. <https://doi.org/10.1332/239868021X16290304343529>
- [223] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 367–385.
- [224] Huixin Tian, Chris Kanich, Jason Polakis, and Sameer Patil. 2020. Tech Pains: Characterizations of Lived Cybersecurity Experiences. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 250–259. <https://doi.org/10.1109/EuroSPW51379.2020.00040>
- [225] Christian Tiefenau, Maximilian Häring, Katharina Krombholz, and Emanuel von Zezschwitz. 2020. Security, Availability, and Multiple Information Sources: Exploring Update Behavior of System Administrators. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. 239–258. <https://www.usenix.org/conference/soups2020/presentation/tiefenau>

- [226] Alan J. Tomassetti, Reeshad S. Dalal, and Seth A. Kaplan. 2016. Is Policy Capturing Really More Resistant Than Traditional Self-Report Techniques to Socially Desirable Responding? *Organizational Research Methods* 19, 2 (April 2016), 255–285. <https://doi.org/10.1177/1094428115627497>
- [227] Andrew Tomlinson, Simon Parkin, and Siraj Ahmed Shaikh. 2022. Drivers and barriers for secure hardware adoption across ecosystem stakeholders. *Journal of Cybersecurity* 8, 1 (Aug. 2022). <https://doi.org/10.1093/cybsec/tyac009>
- [228] Roger Tourangeau, Lance J. Rips, and Kenneth Rasinski. 2000. *The Psychology of Survey Response*. Cambridge University Press.
- [229] UK National Cyber Security Center. 2024. Smart devices: new law helps citizens to choose secure products. <https://www.ncsc.gov.uk/blog-post/smart-devices-law>
- [230] European Union. 2014. Consolidated text: Directive 2014/53/EU of the European Parliament and of the Council of 16 April 2014 on the harmonisation of the laws of the Member States relating to the making available on the market of radio equipment and repealing Directive 1999/5/EC (Text with EEA relevance). *Official Journal of the European Union* (2014). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02014L0053-20241228>
- [231] European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). *Official Journal of the European Union* (2016). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [232] European Union. 2017. Consolidated text: Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC (Text with EEA relevance). *Official Journal of the European Union* (2017). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02017R0745-20250110>
- [233] European Union. 2017. Regulation (EU) 2017/746 of the European Parliament and of the Council of 5 April 2017 on in vitro diagnostic medical devices and repealing Directive 98/79/EC and Commission Decision 2010/227/EU. *Official Journal* (2017). <http://data.europa.eu/eli/reg/2017/746/2023-03-20>
- [234] European Union. 2019. Consolidated text: Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing Regulation (EU) No 526/2013 (Cybersecurity Act) (Text with EEA relevance). *Official Journal of the European*

- Union* (2019). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02019R0881-20250204>
- [235] European Union. 2019. Consolidated text: Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing Regulation (EU) No 526/2013 (Cybersecurity Act) (Text with EEA relevance). *Official Journal of the European Union* (2019). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02019R0881-20250204>
- [236] 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 (Text with EEA relevance.). *Official Journal of the European Union* (2019). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019L0770>
- [237] 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 (Text with EEA relevance.). *Official Journal of the European Union* (2019). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019L0771>
- [238] European Union. 2019. Guidance on Cybersecurity for medical devices. *Official Journal of the European Union* (2019). [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2021\)689333](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)689333)
- [239] European Union. 2021. Questions and Answers: Strengthening cybersecurity of wireless devices and products. https://ec.europa.eu/commission/presscorner/detail/en/QANDA_21_5635
- [240] European Union. 2022. Consolidated text: Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive) (Text with EEA relevance)Text with EEA relevance. *Official Journal of the European Union* (2022). <https://eur-lex.europa.eu/eli/dir/2022/2555/2022-12-27/eng>
- [241] European Union. 2023. Regulation (EU) 2023/988 of the European Parliament and of the Council of 10 May 2023 on general product safety, amending Regulation (EU) No 1025/2012 of the European Parliament and of the Council and Directive (EU) 2020/1828 of the European Parliament and the Council, and repealing Directive 2001/95/EC of the European Parliament and of the Council and Council Directive 87/357/EEC (Text with EEA relevance). *Official Journal of the European Union* (2023). <https://eur-lex.europa.eu/eli/reg/2023/988/oj/eng>

- [242] European Union. 2024. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC (Text with EEA relevance). *Official Journal of the European Union* (2024). <https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng>
- [243] European Union. 2024. Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act) (Text with EEA relevance). *Official Journal of the European Union* (2024). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R2847>
- [244] 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 preprint arXiv:2310.00942* (2023).
- [245] Veerle van Harten, Carlos Hernandez Ganán, Michel van Eeten, and Simon Parkin. 2025. “All Sorts of Other Reasons to Do It”: Explaining the Persistence of Sub-optimal IoT Security Advice. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 387, 19 pages. <https://doi.org/10.1145/3706598.3713719>
- [246] Kami Vaniea and Yasmeen Rashidi. 2016. Tales of Software Updates: The process of updating software. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3215–3226. <https://doi.org/10.1145/2858036.2858303>
- [247] Luis Vargas, Logan Blue, Vanessa Frost, Christopher Patton, Nolen Scaife, Kevin RB Butler, and Patrick Traynor. 2019. Digital Healthcare-Associated Infection: A Case Study on the Security of a Major Multi-Campus Hospital System.. In *NDSS*.
- [248] 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.
- [249] Swaathi Vetrivel. 2025. *The Signals We Send: Analysing the Market Signals for IoT Security and Privacy*. Delft University of Technology.
- [250] 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.

- [251] Hans-Martin von Stockhausen and Marc Rose. 2020. Continuous security patch delivery and risk management for medical devices. In *2020 IEEE International Conference on Software Architecture Companion (ICSA-C)*. IEEE, 204–209.
- [252] Amy Wang. 2018. Nest cam security breach: A hacker took over a baby monitor and broadcast threats, Houston parents say. <https://www.washingtonpost.com/technology/2018/12/20/nest-cam-baby-monitor-hacked-kidnap-threat-came-device-parents-say/>
- [253] Dingding Wang, Muhui Jiang, Rui Chang, Yajin Zhou, Hexiang Wang, Baolei Hou, Lei Wu, and Xiapu Luo. 2024. An Empirical Study on the Insecurity of End-of-Life (EoL) IoT Devices . *IEEE Transactions on Dependable and Secure Computing* 21, 04 (July 2024), 3501–3514. <https://doi.org/10.1109/TDSC.2023.3334017>
- [254] Wang Wei. 2018. Casino Gets Hacked Through Its Internet-Connected Fish Tank Thermometer. <https://thehackernews.com/2018/04/iot-hacking-thermometer.html>
- [255] Charles Weir, Anna Dyson, and Dan Prince. 2023. Do You Speak Cyber? Talking Security With Developers of Health Systems and Devices. *IEEE Security & Privacy* 21, 1 (2023), 27–36. <https://doi.org/10.1109/MSEC.2022.3221616>
- [256] Charles Weir, Ben Hermann, and Sascha Fahl. 2020. From Needs to Actions to Secure Apps? The Effect of Requirements and Developer Practices on App Security. In *29th USENIX Security Symposium (USENIX Security 20)*. 289–305.
- [257] Harald Wieser, Nina Tröger, and Renate Hübner. 2015. The consumers' desired and expected product lifetimes. *Product Lifetimes And The Environment* (2015).
- [258] Markus Willing, Christian Dresen, Eva Gerlitz, Maximilian Haering, Matthew Smith, Carmen Binnewies, Tim Guess, Uwe Haverkamp, and Sebastian Schinzel. 2021. Behavioral responses to a cyber attack in a hospital environment. *Scientific Reports* 11, 1 (Sept. 2021), 19352. <https://doi.org/10.1038/s41598-021-98576-7>
- [259] Daniel W Woods and Aaron Ceros. 2021. Blessed Are The Lawyers, For They Shall Inherit Cybersecurity. In *New Security Paradigms Workshop*. 1–12.
- [260] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019. Defending my castle: A co-design study of privacy mechanisms for smart homes. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–12.
- [261] Omer Yoachimik and Jorge Pacheco. 2025. Record-breaking 5.6 Tbps DDoS attack and global DDoS trends for 2024 Q4. <https://blog.cloudflare.com/ddos-threat-report-for-2024-q4/>
- [262] 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.

- [263] Eric Zeng and Franziska Roesner. 2019. Understanding and Improving Security and Privacy in Multi-User Smart Homes: A Design Exploration and In-Home User Study. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, 159–176.
- [264] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh. 2021. “Did you know this camera tracks your mood?”: Understanding Privacy Expectations and Preferences in the Age of Video Analytics. *Proceedings on Privacy Enhancing Technologies* 2021, 2 (Jan. 2021). <https://doi.org/10.2478/popets-2021-0028>

A

APPENDIX FOR CHAPTER 2

A.1. SURVEY INSTRUMENT

Q1: Device Introduction

In this study we would like to learn about your opinions on internet connected devices. But what exactly are internet connected or 'smart' devices? A growing number of physical things around us are being equipped with electronic capabilities and connected to the internet. They become increasingly smart. Examples of such devices are smartphones, smart speakers with voice assistants, smart security cameras, smart thermostats, smartwatches, or cars with increasing connectivity. All these devices have several features in common:

- They have one or more sensors that collect information around them, like microphones, video cameras, a thermometer, or GPS.
- They are connected to the internet, for example via WiFi or Cellular networks.
- Thus, they can communicate with other devices, mobile applications, or servers (e.g. websites or the 'Cloud').
- They allow features for users which would not be possible if they were not connected.

Q2: Device experience: Do you have personal experience with such devices? Please select all internet connected products you have used at least once during the last four weeks:

Personal Computer (Desktop PC and/ or Laptop) | Smartphone | Tablet | Smart TV | Smart Speaker (e.g. Google Nest, Amazon Echo) | Smart Watch and/ or fitness tracker (e.g. FitBit, Apple Watch) | Gaming Console (e.g. Playstation, Xbox) | Smart Lighting (e.g. Philips Hue, Wyze bulb) | Smart Thermostat | Internet Connected Security Camera | Smart doorlock and/ or doorbell | Robot vacuum | Printer with internet / WiFi connection | Car with at least one internet connected component (e.g. Infotainment, Board computer, WiFi or cellular connectivity) | Other device(s): (- *Open text field* -)

Q3: Vignettes: On the following pages, you will read seven fictional short stories about people and their experiences with internet connected devices. After each story, we will ask you some questions about your impressions and thoughts on the account. Please read the stories carefully and answer the questions as genuinely as possible. There are no right or wrong answers, and we do not want to test your performance in any way. Instead, we would appreciate your honest opinion and impressions. **One exemplary vignette (of seven):** *[text the same as example Vignette in Methodology section in main body of paper, without numbering of vignette elements.]*

Q3.1: If you had to predict, how likely do you think a real manufacturer would respond this way considering the circumstances? *Extremely unlikely, Unlikely, Somewhat unlikely, Neither likely nor unlikely, Somewhat likely, Likely, Extremely likely*

Q3.2: The manufacturer's response to the situation is appropriate considering the circumstances. *Strongly disagree, Disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Agree, Strongly agree*

(If answer to previous question was *Strongly disagree, Disagree, or Somewhat disagree*.) What exactly was inappropriate about the manufacturer's response? (- Open text field -)

Q3.3: Alex's response is a suitable way to move forward from this situation. *Strongly disagree, Disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Agree, Strongly agree*

(If answer to previous question was *Strongly disagree, Disagree, or Somewhat disagree*.) What should Alex do instead? (- Open text field -)

Q3.4: The situation Alex is faced with is concerning.¹ *Strongly disagree, Disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Agree, Strongly agree*

Q3.5: The situation described in the story is realistic. *Strongly disagree, Disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Agree, Strongly agree*

(If answer to previous question was *Strongly disagree, Disagree, or Somewhat disagree*.) What do you think is unrealistic about this story? (- Open text field -)

Q4: Confidence in previous answers: How confident do you feel about your answers to the previous stories? *Very confident, Somewhat confident, Somewhat unconfident, Very unconfident, I don't know*

(If answer to previous question was *Somewhat unconfident* or *Very unconfident*.) What exactly makes you feel unconfident about your answers? (- Open text field -)

Q5: Personal experiences: Did the previous stories remind you of any personal experiences you have had with electronic devices? *Yes, No, I don't know* (If answer to previous question was *Yes*.) Please briefly tell us about your personal experiences you were reminded of: (- Open text field -)

Q6: Attention check: We would like to learn what your favorite device is. Please select 'Washing machine' from the list below. This is an attention check. Based on the text you read above, what device have you been asked to enter? *Smartphone, Washing machine, Laptop, Voice assistant*

Q7: Age: How old are you? *Under 18, 18-24 years old, 25-34 years old, 35-44 years old, 45-54 years old, 55-64 years old, 65+ years old, Prefer not to say*

Q8: Gender: Which gender do you identify with? *Male, Female, Non-binary / third gender, Prefer not to say*

Q9: Country: In which country do you currently reside? (*Select from list*)

Q10: Employment status: Which option best describes your current status? *Employed, Unemployed, Student, Retired, Other* (If answer to previous question was *Employed*.) Which field are you working in? (*Select from list*)

Q11: Last comments: Would you like to make any last comments or remarks about this survey? (- Open text field -)

Q12: Debriefing: Thank you for your responses! This study was about security and privacy issues of internet connected devices. We did not disclose this at the beginning of the survey directly to not influence your responses in any way. Security and privacy is a sensitive topic which becomes increasingly important as the number of different smart devices in our environment grows rapidly. In this survey, we wanted to measure your expectations about responses of different actors for varying security and privacy issues and smart devices. By reading and responding to different short stories, you helped us a lot to do so.

¹We did not focus on this scale in our analysis due to space limitations and no important relationships over what the other scales already showed.

A.2. REGRESSION TABLE

Coefficient	Appropriateness ratings						Likelihood ratings						User suitability ratings					
	Model 1: Security		Model 2: Privacy		Model 3: Security		Model 4: Privacy		Model 5: Security		Model 6: Privacy		Model 7: Security		Model 8: Privacy			
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error		
Intercept	1.65 ***	0.14	2.61 ***	0.12	3.80 ***	0.14	4.52 ***	0.08	3.15 ***	0.15	4.63 ***	0.11						
Device:																		
Connected car	0.30	0.17	-0.03	0.14	0.27	0.19	0.05	0.08	-0.35	0.18	-0.34 ***	0.08						
Smart speaker	0.05	0.16	-0.07	0.15	0.19	0.18	0.20 *	0.08	0.06	0.17	-0.12	0.08						
Smart washing machine	0.64 ***	0.17	-0.08	0.14	0.54 ***	0.18	0.12	0.08	-0.26	0.18	-0.03	0.08						
Smartphone	0.39 *	0.17	-0.07	0.15	0.15 *	0.19	0.21 **	0.08	0.07	0.18	-0.22 **	0.08						
Smartwatch	0.32	0.16	-0.04	0.15	0.28	0.18	0.13	0.08	0.36 *	0.18	-0.06	0.08						
Security/ Privacy Event:																		
IoT ransomware	-0.10	0.06			-0.02	0.07			0.16 **	0.06								
Unauthorized data access	-0.01	0.06			0.05	0.07			0.08	0.06								
DDoS (reference)																		
Forced data collection			0.19	0.15			0.16 *	0.06			0.47 ***	0.14						
Third party sharing			0.44 **	0.15			0.15 **	0.06				0.13						
No consent (reference)																		
Manufacturer response:																		
(Sec) Inform users	2.37 ***	0.19			1.02 ***	0.19			0.30	0.18								
(Sec) Recall	4.21 ***	0.19			1.14 ***	0.19			-0.17	0.17								
(Sec) Patch	3.54 ***	0.19			1.89 ***	0.18			0.45 *	0.18								
(Sec) No response (reference)																		
(Priv) Inform via privacy policy			2.14 ***	0.06			0.58 ***	0.06			0.34 ***	0.07						
(Priv) Privacy S-W update			2.12 ***	0.06			0.47 ***	0.05			0.18 **	0.06						
(Priv) No response (reference)																		
User response:																		
Advice online forums									2.01 ***	0.09								
Return device									2.65 ***	0.09								
Tech. mitigation									1.73 ***	0.09								
Turn off									1.95 ***	0.09								
Keep using (reference)																		
Interaction Effects:																		
ManuResp*Device:	0.602		Device*PrivEvent:		ManuResp*Device:		None found		ManuResp*Device:		Event*UserResp:							
Inform*Car (-0.75 **)		Car*Forced (0.43 **)	Car*Forced (0.43 **)		Recall*WashMachi (-0.89 **)				Inform*Smartwatch (-0.81 **)		Forced*Return (-0.82 **)							
**WashMachi (-0.51*)		WashMachi*Forced (0.78 **)	WashMachi*Forced (0.78 **)		Patch*WashMachi (-0.69 *)				Recall*Car (0.54 *)		Forced*TurnOff (-0.98 **)							
**Smartphone (-0.65 **)		Smartphone*Forced (0.57 **)	Smartphone*Forced (0.57 **)						Patch*Smartphone (-0.98 **)									
Recall*WashMachi(-0.63 **)		Smartwatch*Forced (0.65 **)	Smartwatch*Forced (0.65 **)															
**Smartphone (-0.87 **)																		
Conditional R ² :	0.602		0.464		0.217		0.248		0.306		0.172							
ICC:	0.15		0.18		0.09		0.22		0.05		0.11							

* p<0.05 ** p<0.01 *** p<0.001

Table A.1: Regression models predicting 1) Appropriateness of manufacturer response, 2) Likelihood of manufacturer response, and 3) Suitableness of user response with the vignette factors; Device, Security/ Privacy Event, Manufacturer response, and User response. Each model was run separately for security and for privacy. Scale of measurement: For Model 1, 2, 5, and 6: 7-point Likert scale with 1 = ‘Strongly disagree’ to the statement ‘The manufacturer’s response to the situation is appropriate considering the circumstances.’ or ‘Alex’s response is a suitable way to move forward from this situation.’ For Model 3 and 4: 7-point Likert scale with 1 = ‘Extremely unlikely’ and 7 = ‘Extremely likely’ to the statement ‘If you had to predict, how likely do you think a real manufacturer would respond this way considering the circumstances?’ For Model 1, there were slight regional differences, where participants from Central and South America rated manufacturer responses significantly more appropriate on average than participants from North America (Coefficient = 0.45, SE = 0.20, p<0.05). Other participant characteristics (Age, Gender, Count of used IT device) did not show an effect. We only present statistically significant interaction terms here for space considerations.

B

APPENDIX FOR CHAPTER 3

B.1. SURVEY INSTRUMENT

Survey description as advertised on Prolific: In this survey, we would like to learn about how you use smart devices and your general thoughts about such products. It will take approximately 18 minutes to complete. We will ask you some questions about any smart devices you are currently using and show you several descriptions of different smart devices to learn about your thoughts about them. These questions will not test your performance in any way, as we are only interested in your perspective. The survey will also ask for basic demographic information like age, gender, occupation, and country of residence.

Q1: Introduction In this survey, we want to learn your opinions on smart devices. There will be no right or wrong answers - we are only interested in your experiences and your opinion of these products. But what exactly are smart devices? A growing number of physical things around us are being equipped with electronic capabilities and connected to the internet and apps. Examples of such devices are smart speakers with voice assistants, smart security cameras, smart thermostats, or smartwatches. These devices have several features in common: They have sensors that collect information around them, like microphones, video cameras, a thermometer, or GPS. They can be connected to the internet, for example, via WiFi or cellular networks. Thus, they can communicate with other devices, mobile apps, or services. They provide features that would not be possible if they were not connected.

Q2: Device usage Which of the following types of smart devices do you own and use? Please select from the list below.

Smartphone, Smartwatch and/ or fitness tracker, Smart TV, Smart speaker with voice assistant (e.g., Amazon Echo or Google Nest), Media streaming device (e.g., Chromecast or music streamer), Smart lighting (smart light bulbs, light switches), Smart hub, Robot vacuum cleaner, Smart washing machine, Smart fridge, Smart baby monitor, Internet connected security camera (indoor and / or outdoor), Smart video doorbell, Smart door-lock, Smart sensor (e.g., smoke detector, motion detector, air quality sensor), Smart thermostat, Wi-Fi solar inverter, Printer with internet / WiFi con-

nection, Internet router, Other(s); None of the above

Q2.1: Device model/brand We selected some of the smart device types you previously indicated to own. For each of them, please indicate the specific device's brand and/or model you are currently using. If you are using multiple devices of the same type (e.g., several different cameras), please pick one and focus your answers on that device.

Brand and/or model name, I don't know (For a list of three randomly chosen devices from Q2)

Q2.2: Start usage When answering the following questions, please focus on these specific smart device models you just described (e.g., the exact smartphone you are currently using), rather than the general device type (e.g., smartphones in general).

For the particular device you are currently using, when did you start using it, as far as you can recall? Please select the year.

2024, 2023, 2022, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014 or earlier, I don't know (For the list of the previous devices)

Q2.3: Condition For the particular device you are currently using, was it new or previously owned (second-hand) when you started using it? *New, Previously owned, I don't know* (For the list of the previous devices)

Q2.4: Purchase channel Where was the device you are currently using purchased from? *Online store, Physical store, I don't know, It wasn't purchased* (For the list of the previous devices)

Q2.5: Shop What is the name of the store where the device was purchased?

Shop name, I don't know (For the list of the previous devices)

Q2.6: Future use Looking ahead, for approximately how many more years do you plan to keep using the particular device you are currently using?

Less than a year, 1 year, 2 years, 3 years, 4 years, 5 years, 6 - 10 years, More than 10 years, I don't know (For the list of the previous devices)

Q2.7: Motivation for planned end of use You indicated that you are planning to not use some of the devices for much longer. Why do you plan to stop using them soon? Please select the choices that best describe your reasons.

Experiencing issues and flaws with it, Desire for

a newer device, Dissatisfied with brand or manufacturer, Concerns about my data, Concerns about cybersecurity, Software update support ended, Incompatible with other devices and/or services, I don't know, Other reasons (For the previous devices planned to be used for less than a year)

Q2.8: Update status Does the particular device you are currently using still receive software updates, as far as you know? (This refers to updates to the device itself, not any companion app or third-party apps running on the device.) *Yes, it still receives software updates, No, it no longer receives software updates, I don't know if it still receives software updates* (For the list of the previous devices)

Q3: Attention check Please select 'Agree' for both Smart Washing Machine and Smartwatch. This is an attention check.

Q4: Recently stopped devices Now we would like to learn about any smart devices that you have recently stopped using, for instance to replace them.

Q4.1: Devices Have you recently stopped using any smart devices? Please select the type of the smart device you stopped using. If you haven't, just select "None of the above". *Smartphone, Smartwatch and/or fitness tracker, Smart TV, Smart speaker with voice assistant (e.g., Amazon Echo or Google Nest), Media streaming device (e.g., Chromecast or music streamer), Smart lighting (smart light bulbs, light switches), Smart hub, Robot vacuum cleaner, Smart washing machine, Smart fridge, Smart baby monitor, Internet connected security camera (indoor and / or outdoor), Smart video doorbell, Smart door-lock, Smart sensor (e.g., smoke detector, motion detector, air quality sensor), Smart thermostat, Wi-Fi solar inverter, Printer with internet / WiFi connection, Internet router, Other(s);, None of the above*

Q4.2: Past use duration For approximately how many years did you use the device in total? If you stopped using more than one device of the same type (e.g., several different cameras), please pick one and focus your answers on that device.

Less than a year, 1 year, 2 years, 3 years, 4 years, 5 years, 6 - 10 years, More than 10 years, I don't know (For the list of three randomly chosen de-

vices from Q4.1)

Q4.3: Action with old device What did you do with the device once you stopped using it? Please select the option that best describes your action.

Sold it, Returned it, Discarded it, Traded it in, Gifted it, Still have it but not in use, I don't know, Something else (For the list of the previous devices)

Q4.4: Motivation Why did you stop using the device? Please select the choice(s) that best describe your reasons. Multiple answers are possible.

It stopped working, Experiencing issues and flaws with it, Desire for a newer device, Dissatisfied with brand or manufacturer, Concerns about my data, Concerns about cybersecurity, Software update support ended, Incompatible with other devices and/or services, I don't know, Other reasons (For the list of the previous devices)

Q5: Device vignettes In the following, you will read short descriptions of three different conventional and smart devices. For each one, we will ask you how long you expect such a device to last. When answering, please imagine you are buying such a device today. Realistically, for how many years would you expect it to last under normal intensity of use until it cannot be used for its intended purpose anymore?

Q5.1: Vignette conventional device
A washing machine. It has different washing programs for different textiles that the user can select on the machine's buttons

Q5.1.1: Expected lifetime If you had to predict, for how many years do you expect such a device to last? Please write your answer as a number in the field below.

(Open text field for integer text input)

Q5.1.2: Factors considered What aspects did you take into account when estimating the number of years? *(Open text field)*

The other devices vignettes:
A home printer. It can print documents in different formats and colors and can be connected to a computer with a cable to receive print jobs.

A vacuum cleaner. It can be used to clean the home's floor and has different suction strengths for different surfaces, which can be selected on the machine's buttons.

A thermostat. It is installed in the home's heating system and can be used to adjust the home's temperature on the thermostat's controls to set the desired temperature.

A smoke detector. It can continuously monitor the home's air for smoke and potential fire hazards and alert occupants with an audible alarm.

A solar inverter. It is connected to a home's solar panels and converts the energy from the solar panels into electricity for the home.

A smart washing machine. It has different washing programs for different textiles that the user can select on the machine's buttons. It can be connected to the internet via the home Wi-Fi, so the user can check the status of their laundry and control the machine remotely through a companion app on their mobile phone.

An internet-connected home printer. It can print documents in different formats and colors and can be connected to a computer with a cable or wirelessly to receive print jobs. It can be connected to the internet via the home Wi-Fi, so the user can send documents to print and check the printer's status remotely.

A robot vacuum cleaner. It can map and clean the home's floor on its own. It is connected to the internet via the home Wi-Fi, so the user can monitor its cleaning progress, set schedules, and control it remotely through a companion app on their mobile phone.

A smart thermostat. It is installed in the home's heating system and can be used to adjust the home's temperature on the thermostat's controls to set the desired temperature. It can be connected to the internet via the home Wi-Fi, so the user can check and control the temperature settings remotely through a companion app on their mobile phone.

A smart smoke detector. It can continuously monitor the home's air for smoke and potential fire hazards and alert occupants with an audible alarm. It can be connected to the internet via the home Wi-Fi, so the user can receive alerts and check its status remotely through a companion app on their mobile phone.

A Wi-Fi solar inverter. It is connected to a home's solar panels and converts the energy from the solar panels into electricity for the home. It can be connected to the internet via the home

Wi-Fi, so the user can track energy production and monitor system performance remotely through a companion app on their mobile phone.

A home router. It keeps up and manages the home Wi-Fi and connects the devices in the home to the internet via a broadband connection. The user can manage connected devices, monitor the activity in the Wi-Fi network, and adjust settings through a web portal.

A smart security camera. It can continuously collect video recordings of the home or its surroundings. It is connected to the internet via the home Wi-Fi, so the user can check the video feed remotely through a companion app on their mobile phone.

A smart door lock. It is installed at the home's entrance door to lock and unlock it. It can be connected to the internet via the home Wi-Fi, so the user can remotely lock or unlock the door, monitor entry logs, or grant or revoke access to guests through a companion app on their mobile phone.

A smartwatch. It can continually track various health metrics of the user, be used for navigation, receiving messages, and can have various apps installed on it. It can be connected to the internet via a cellular network, so the user can use such features directly on the watch or through a companion app on their mobile phone.

A smart speaker. It can continually listen for voice commands to perform features like streaming music, setting reminders, getting weather updates, or controlling other connected devices in the home. It is connected to the internet via the home Wi-Fi, so the user can use these features via the smart speaker or manage the speaker via a companion app on their mobile phone.

Q6: Attention check We would like to learn what your favorite device is. Please select Voice assistant from the list below. This is an attention check.

Based on the text you read above, what device have you been asked to enter? *Smartphone, Washing machine –Laptop, Voice assistant, Smart Watch*

Q7: End of update support expectations Many smart devices receive software updates for a limited time after purchase. Once a device no longer receives software updates, how do you expect this will impact its use? Please select all that ap-

ply.

Device's performance declines, Some critical features stop working, affecting important functionality, Some non-essential features stop working, affecting 'nice-to-have' aspects, New features are no longer added, Less compatibility with other devices or services, Customer support services are no longer available, Cybersecurity weaknesses in the device are no longer fixed, Potential cybersecurity risks to the device are no longer tracked (new threats may go unnoticed), Other, namely, I don't expect any changes, I'm not sure

Q7.1: Own experience Have you experienced any of these changes with your own smart devices after software update support ended? Yes, No, I don't know (For the selected answers in Q7)

Q8: Update support duration Coming back to the two smart devices you have previously been presented with. For how many years do you expect such a smart device to receive software updates after purchase, if at all? Please write your answer as a number in the left-most column, or select one of the other options. *Number of years: 0, I don't know, I don't expect it to receive updates* (For the two smart devices seen in the vignettes.)

Q8.1: End of update support expectations You indicated that software updates for the smart device(s) will likely stop before the device's actual end of life. In what ways, if any, do you expect the device to work differently after software updates stop and it is still used? *(Open text field)* (Only shown if response to Q5.1.1 (expected lifetime) was lower than Q8 (expected update support))

Q9: Security considerations Since smart devices can be connected to the internet, they can be exposed to cybersecurity risks. Cybersecurity for smart devices focuses on protecting the devices, users, and their data from potential attacks through online connectivity.

Q9.1: Concerns Do you have any cybersecurity concerns regarding the use of your own smart devices? Please select all that apply, if any. *My personal data being accessed by unauthorized individuals., My devices being hacked and used maliciously without my knowledge., My devices being damaged or made unusable by a cyberat-*

tack., Insufficient customer support in case of cybersecurity issues., My personal data being shared with third parties by the manufacturer without my consent., More personal data collected from the device than I expect., Other:, I'm not concerned., I'm not sure.

Q9.2: Mitigation actions Do you take any actions to address your cybersecurity concerns with your smart devices? Please select all that apply, if any. *Ensure the latest software updates are installed., Replace devices once software update support has ended., Disconnect devices from the network when not in use., Change device settings, such as passwords., Check online for news or warnings about my device., Be cautious about how I use and behave around my device., Choose brands that I trust when buying devices., Others:, None of the above.*

Q9.3: Prolonged use Please rate your agreement with the following statement: I would use my smart devices for longer if they continued to receive software updates that fix cybersecurity weaknesses.

Strongly agree, Agree, Neither agree nor disagree, Disagree, Strongly disagree

Q10: Demographics Please answer some final questions about yourself to complete the survey.

Q10.1: Age How old are you? *Under 18, 18-24 years old, 25-34 years old, 35-44 years old, 45-54 years old, 55-64 years old, 65+ years old, Prefer not to say*

Q10.2: Gender Which gender do you identify with? *Male, Female, Non-binary / third gender, Prefer not to say*

Q10.3: Country In which country do you currently reside? *(List)*

Q10.4: Employment Which option best describes your current status? *Employed / self-employed, Unemployed, Student, Retired, Other*

Q10.5: Employment field Which option best describes the field you are currently working in? *Business or sales, Administrative support, Research, Hospitality service – Education, Computer engineer or IT, Engineer in other fields, Medical, Le-*

gal, Homemaker, Skilled labor, Art or writing, Government, Other, Prefer not to say

Q11: Debriefing Thank you for your responses!

This study was about your expectations of smart device's lifetime and the role that security and privacy play. We did not disclose this explicitly at the beginning of the survey to not influence your responses in any way. Security and privacy is a sensitive topic which becomes increasingly important as the number of different smart devices in our environment grows rapidly. At the same time, regulations increasingly demand a longer time for these devices to receive software updates to protect them from attacks. In this survey, we wanted to measure your expectations about how long such devices will last, as this could affect how long manufacturers will provide such security updates. By reading and responding to the different descriptions of smart devices, you helped us a lot to research this topic.

B.2. REGRESSION TABLE

	Coefficient	Standard error
Intercept	6.431***	0.054
Device Type (Ref: Smart speaker)		
Conventional printer	0.072	0.052
Conventional smoke detector	0.211***	0.053
Conventional solar inverter	0.569***	0.053
Conventional thermostat	0.527***	0.053
Conventional vacuum cleaner	0.054	0.053
Conventional washing machine	0.301***	0.052
Smart security camera	0.061	0.043
Robot vacuum cleaner	-0.126**	0.043
Home router	0.072	0.043
Smart door lock	0.116*	0.043
Connected printer	0.091*	0.043
Smart smoke detector	0.244***	0.043
Smart thermostat	0.434***	0.043
Smart washing machine	0.232***	0.043
Smartwatch	-0.221***	0.043
WiFi solar inverter	0.469***	0.043
Country (Ref: Spain)		
CountryFrance	-0.041	0.046
CountryGermany	0.092	0.046
CountryNetherlands	0.086	0.046
CountryPoland	-0.141***	0.046
Device Experience		
Device category experience	0.080**	0.024
Country:Smartness (Ref: Poland)		
Spain:Smartness	-0.049	0.045
France:Smartness	-0.071	0.045
Germany:Smartness	-0.085*	0.045
Netherlands:Smartness	-0.122**	0.045
AIC	3477.401	
BIC	3645.336	
Num. observations	2974	
Num. groups: Participant	993	
Conditional R^2	0.464	
Marginal R^2	0.166	

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table B.1: Regression model predicting expected device lifetime in the vignettes with Device Type, Country, respondents' Experience with the Device Category, and an Interaction between Country and the device's Smartness. "Ref" = Reference level. The outcome variable was log-transformed, and coefficients were exponentiated to denote an average increase or decrease in expected device lifetime, holding all other factors constant. Standard errors are not exponentiated.

B.3. OPEN TEXT RESPONSE CODES

Device	Factor	Count	Proportion
Washing Machine (Conventional)	Device type	47	19.03%
	Usage	43	17.41%
	Own experience	38	15.38%
	Components	22	8.91%
	Device quality	16	6.48%
Printer (Conventional)	Usage	63	25.82%
	Own experience	55	22.54%
	Device type	26	10.66%
	Components	23	9.43%
	Device quality	14	5.74%
Vacuum cleaner (Conventional)	Usage	53	24.65%
	Own experience	33	15.35%
	Device type	27	12.56%
	Components	20	9.30%
	Device quality	17	7.91%
Thermostat (Conventional)	Usage	58	24.79%
	Device type	41	17.52%
	Complexity	19	8.12%
	Own experience	17	7.26%
	Components	15	6.41%
Smoke detector (Conventional)	Usage	60	26.32%
	Complexity	38	16.67%
	Device type	23	10.09%
	Components	20	8.77%
	Safety	18	7.89%
Solar inverter (Conventional)	Device type	33	13.98%
	Price	32	13.56%
	Usage	29	12.29%
	Other / similar device	28	11.86%
	Estimate	17	7.20%
Smart washing machine (Smart)	Device type	70	26.32%
	Usage	33	12.41%
	Own experience	27	10.15%
	Smart complexity	22	8.27%
	Better alternative	14	5.26%
Connected printer (Smart)	Usage	56	21.71%
	Own experience	41	15.89%
	Device type	39	15.12%
	Updates	19	7.36%
	Components	18	6.98%

(Continued on next page.)

Table B.2: Factors taken into account by respondents when estimating devices' lifetimes. Factors were derived by coding open-text responses to the question: "What aspects did you take into account when estimating the number of years?". Only the five most prevalent factors per device are listed.

(Continued from previous page.)

Device	Factor	Count	Proportion
Robot vacuum <i>(Smart)</i>	Usage	68	25.66%
	Device type	30	11.32%
	Environment	22	8.30%
	Components	20	7.55%
	Own experience	19	7.17%
Smart thermostat <i>(Smart)</i>	Usage	48	19.59%
	Device type	33	13.47%
	Updates	29	11.84%
	Complexity	20	8.16%
	Estimate	15	6.12%
Smart smoke detector <i>(Smart)</i>	Usage	42	17.87%
	Device type	39	16.60%
	Complexity	19	8.09%
	Safety	19	8.09%
	Smart complexity	15	6.38%
WiFi Solar inverter <i>(Smart)</i>	Estimate	31	12.60%
	Device type	28	11.38%
	Usage	28	11.38%
	Price	22	8.94%
	Other / similar device	20	8.13%
Home router <i>(Smart)</i>	Own experience	49	19.92%
	Usage	40	16.26%
	Better alternative	38	15.45%
	(In-)compatibility	29	11.79%
	Device type	23	9.35%
Smart security camera <i>(Smart)</i>	Usage	37	14.12%
	Environment	32	12.21%
	Updates	27	10.31%
	Better alternative	25	9.54%
	Device quality	23	8.78%
Smart doorlock <i>(Smart)</i>	Usage	39	16.05%
	Device type	33	13.58%
	Estimate	18	7.41%
	Updates	18	7.41%
	Better alternative	15	6.17%
Smartwatch <i>(Smart)</i>	Usage	51	18.02%
	Better alternative	41	14.49%
	Updates	38	13.43%
	Components	29	10.25%
	Obsolescence	21	7.42%
Smart speaker <i>(Smart)</i>	Usage	38	15.51%
	Device type	34	13.88%
	Updates	32	13.06%
	Better alternative	27	11.02%
	Obsolescence	19	7.76%

B.4. CODEBOOK

Theme	Code	Code description
Static device quality	Device Type	The general sort of device was considered (e.g., "Washing machines last long")
	Conventional / Smart Device type	Lifetime estimate was explicitly based on the conventional counterpart of the smart device (e.g., "A normal vacuum cleaner lasts that long.")
Dynamic device quality	Device Quality Components	The device build quality would be taken into account. Components of the device and their respective longevity were mentioned (e.g., "The battery will degrade quickly.")
	Manufacturer disclosure	What is disclosed about the device via marketing, the packaging, manuals, or the website.
	Brand	The device's brand would be considered.
	Similar / Other Device	A similar device was taken as an heuristic (e.g., solar panels as an heuristic for solar inverters, or smartphones for smart watches).
	Security	Security factors were considered (e.g., new encryption protocols or vulnerabilities).
	Safety	The safety relevance of the device was highlighted, requiring it to be reliable (e.g., reliance on a smoke detector)
	Modern	Modern technology in general was considered (e.g., "Modern devices break earlier" or "Modern devices are more reliable").
	Price	The device's purchase price was taken into account.
	Country	The device's country of manufacture was considered.
	Complexity	A general perception of the devices build-up complexity (e.g., "It is a complex device, so there are more parts that can break.")
Device use	Smart complexity	It is highlighted explicitly that the added smartness (complexity) has an effect on the device's lifespan.
	Obsolescence	Planned obsolescence was mentioned as a consideration.
	Updates	Updates were explicitly mentioned as a factor.
	Innovation: Better alternatives	Newer, more attractive products available on the market, so that fast innovation drives faster replacement.
	(In-)compatibility	(In-)compatibility with other devices or services that are necessary to run the device.
	Spare parts / Component availability	For how long important components or spare parts are available.
	Upkeep costs	Costs of running the device were considered, like electricity costs or printer ink.
	Usage	The frequency and intensity the device is used by its owners was considered.
	Environment	Environmental factors affecting device lifespan, like dog hairs for a vacuum cleaner or outdoor conditions.
	Maintenance / Installation	Device maintenance and duty of care were considered, or how the device is installed.
Subjective factors	Own experience	Respondents' own experiences with this device.
	Estimate	Respondents emphasized that they estimated the device's lifespan, e.g., because they were not completely sure.
Support structure	Preference	Respondents explicitly mentioned that they considered their normative preference for the lifetime (e.g., "wishful thinking").
	Warranty	The warranty period of the device type was mentioned.
	Regulations	Local regulations were taken into account (e.g., Laws on regular smoke detector replacement).

Table B.3: Final codebook of analysis of open text responses to the question: "What aspects did you take into account when estimating the number of years?"



C

APPENDIX FOR CHAPTER 4

C.1. INTERVIEW PROTOCOL – HDOs

PARTICIPANT BACKGROUND

1. What is your role here at (*organisation*)?
 - (a) For how many years are you already working in this position?
2. What is your role here with respect to medical devices?
 - (a) What type of medical device classes are you responsible for?
3. How many medical devices are there in your organization?
4. How many medical devices in your organization are connected (i.e., to networks)?
 - (a) (If applicable to role) And how many connected medical devices in total are you responsible for?
5. How do you keep track of and manage these connected medical devices?
 - (a) Does this work out well from your perspective?
 - (b) (*If it's not perfect*; What would help you to improve the management of connected medical devices?

PATCHING PROCESS

1. What percentage of the software updates you install on connected medical devices would you estimate are **security** updates? (i.e., they close vulnerabilities in some software or hardware component)
2. Let's say a security update has to be installed on a connected medical device. Can you walk me through the process of how this updating works at your organization?
3. (*If questions 5 - 8 have already been answered before, omit accordingly.*)
4. How is it decided which devices need to be updated and when, and who decides?
5. How do you learn about available updates?
6. Who do you typically interact with during this process? What is their role in this?
7. Are updates tested before they are installed, and if yes, how?
8. Do you have any certifications to be able to install updates on connected medical devices?
9. How often do you have to be (re-)certified?
10. Now I would like to ask you some questions about the last three times you installed a security update on a connected medical devices. What were the last three medical devices you were involved in patching?
11. For each of them, can you tell us...:
 - (a) What kind of medical device was it?
 - (b) From which manufacturer is the device?
 - (c) When did you install the update on this device?
 - (d) How did you learn about the need to update?
 - (e) How often do you usually install updates on this device?
 - (f) How are updates installed on this device? (e.g. remote? manual? by manufacturer technician? Installation steps?)
 - (g) How is the device connected to the rest of the organization?
 - (h) What kind of issue did the security update fix, if you know?
 - (i) What was the time span between the *release* of the update to the *installation*, if you know?
12. What are the biggest challenges you face when it comes to keeping connected medical devices up-to-date?

13. How do you handle these challenges?
14. *If the following topics have not been mentioned previously:*
 - (a) How do you handle connected medical devices that don't receive software updates anymore and stay in service?
 - (b) How do you deal with downtime, when the device is updated and not available for medical use?
 - (c) Were there situations when the installation of an update led to problems with the device's performance? If yes, what do you do in such cases?
15. Are there any security risks you are concerned about when it comes to connected medical devices at your organization?

REGULATORY FACTORS

1. Which national and international regulations do you follow when dealing with connected medical devices, if any?
2. How do these regulations effect the software updating process within your organization, if at all?

HOSPITALS' STANCE TOWARDS MANUFACTURERS

1. How is the manufacturer of the medical device involved in the updating and patching process of medical devices?
2. In your experience, do medical device manufacturers meet their responsibilities in keeping connected medical devices secure at your organization?

CLOSING QUESTION:

1. Thank you very much for taking the time to talk to us. Do you have any closing remarks or things you would like to mention?
2. Can you recommend a colleague at your organization or at another one, who works in a similar position (is involved in patching) and could participate in our study?

C.2. INTERVIEW PROTOCOL – MANUFACTURERS

PARTICIPANT BACKGROUND

1. What is your role here at (*organisation*)?
 - (a) For how many years are you already working in this position?
2. What kind of medical device products are you dealing with in your work?

PATCHING MEDICAL DEVICES

1. How often do you release security updates for your connected medical device products?
 - (a) (*Potential follow-up questions*;)
 - (b) Is this different for different devices? Why?
 - (c) What percentage of the software updates you release for your connected medical device products would you estimate are **security** updates?
 - (d) Are security updates bundled with other kinds of software updates, such as performance or feature updates? If yes, why?
2. How is it decided if and when you release a security update for your connected medical device products?
 - (a) (*Potential follow-up questions*;)
 - (b) Which factors do you consider in this decision?
 - (c) Which actors are involved in this decision? (e.g., (Software) vendors, authorities, customers)?
 - (d) What kind of risks do you want to protect your products against?
 - (e) How are such risks assessed?
 - (f) How do regulations affect the decision if and when to patch?
 - (g) In your experience, what are the biggest challenges when it comes to risk assessment and deciding about security patching?
3. Once it is decided to release a security update for your connected medical device products, how does this update reach the devices at your customers' sites?
 - (a) (*Potential follow-up questions*;)
 - (b) What are the different processes to do so?
 - (c) What are the different processes' prevalence? i.e., Which patching process is most and least common?
 - (d) To what degree can customers decide on how to install security updates?
 - (e) How are customers notified about available security updates?
 - (f) If technicians are sent to install security updates, how often does this happen?
 - (g) Are there differences between countries how this process works?
 - (h) What are you observing how these patching processes work out in practice with your customers?
 - (i) In your experience, what are the biggest challenges when it comes to deploying security updates to medical devices at your customers' sites?
4. Do you keep track of your connected medical device products in circulation and their software versions? If yes, how?
 - (a) (*Potential follow-up questions*;)
 - (b) If you cannot keep track, how do you go about updating them?
 - (c) Do you have means to know if the security update has been installed or not?
 - (d) In your experience, what are the biggest challenges when it comes to tracking the software versions and updates of your medical devices at your customers' sites?

CLOSING QUESTION:

1. Thank you very much for taking the time to talk to us. Do you have any closing remarks or things you would like to mention?

C.3. DEVICE PATCHING INSTANCES

DID	Type	Vendor	HDO	Update frequency	Connections	Install method
1	Docking stations for syringe pumps	A	1	So far, once	To Nurse call system	Via flash card
2	Patient monitoring system	B	1	Upgrade every 2 years	To PDMS, telephones, nurse call system, HIS	Prepare update on server, install on monitors via UI
3	Medical ventilators	C	1	Every 3 - 4 years	To PDMS via server	Not specified
4	Ultrasound System	D	1	Approx. once per year	Doesn't know	Via USB stick
5	Ultrasound System	D	1	Approx. once per year	Doesn't know	Via USB stick
6	Ultrasound System	D	1	Approx. once per year	Doesn't know	Via USB stick
7	Patient monitoring system	B	2	Check for available updates in random interval	Bedside monitors to central stations to HIS. Within own VLAN. Gateway server allows remote updates to central stations.	Prepare update on server, install on monitors via UI
8	Radiology bucky system	E	4	Every 4 - 6 months	To HIS, RIS, and quality system.	Via USB stick
9	MRI	D	4	Every 3 - 4 months	To PACS, RIS, and HIS.	Via USB stick
10	EEG	F	4	Approx. once per year	Doesn't know	Via device UI
11	Laptop as part of medical device	NA	4	Windows 10 / 11 patch frequency	Laptop connected to ECG and ABM.	Via laptop UI
12	MRI	B	5	Approx. once per year	LAN, not in domain	Via device UI
13	Ophthalmic retinal camera	G	6	NA	To PACS via server	Server and devices in tandem (via Laptop)
14	Patient monitoring system	H	6	<i>"Not often"</i> .	Bedside monitors to central stations to EPR via gateway servers. Within own VLAN	Server and devices (via USB stick)
15	Bucky X-Ray system	E	4	Approx. once per year	To PACS and HIS / RIS via LAN	Receive file from vendor, install via device UI
16	MRI	D	4	Approx. once per year	To PACS and HIS / RIS via LAN, remote connection to manufacturer.	Via USB stick
17	Database server connected to sleep monitor	NA	4	Once per month	To EEG devices via workstation PCs, to PACS	On server (remotely)
18	Patient monitoring system	B	7	Every 2 months	Bedside monitors to central stations to HIS. Within own VLAN. Gateway server allows remote updates to central stations.	Prepare update on server, install on monitors via UI
19	Infusion Pumps	I	7	Approx. once per year	To smart docking stations to HIS.	Via laptop connection to device
20	Ventilator	J	7	Only in case of patient risk. <i>"Rarely"</i>	Standalone (currently not connected)	Via laptop connection to device

(Continued on next page.)

DID	Type	Vendor	HDO	Update frequency	Connections	Install method
<i>(Continued from previous page.)</i>						
21	Vessel-sealing coagulation system	K	6	First time	Standalone (currently not connected)	Via laptop connection to device
22	Automated external defibrillator (AED)	L	6	Every 2 years	To HDO-wide AED alert server via WiFi.	Via USB stick
23	Infusion pumps with docking stations	L	6	First time	Docking stations to HIS. Within own VLAN	Via laptop connection to device
24	Diathermy machine	M	8	Every 2 years	Doesn't know	Via laptop connection to device, VPN connection to manufacturer
25	ECG machine	N	9	Every 2 years	Via WiFi to middle-ware to EMR	Via USB stick

Table C.1: List of update cases for different connected medical devices. *DID* = Device ID. Not all cases were (pure) security updates due to common bundling. Nonetheless, the process is the same for security updates. Vendor names were obfuscated. *ABM* = Ambulatory blood pressure monitor. *ECG* = Electrocardiograph. *EEG* = Electroencephalograph. *EMR* = Electronic medical record system. *EPR* = Electronic patient record. *HIS* = Hospital information system. *MRI* = Magnetic resonance imaging device. *PACS* = Picture archiving and communication system. *PDMS* = Patient data management system. *RIS* = Radiology information system.

C.4. CODEBOOKS

CODEBOOK: HEALTHCARE DELIVERY ORGANIZATION

<p>HDO infrastructure</p> <ul style="list-style-type: none"> Departmental structure HDO size Inventory management Network infrastructure Procurement process Regulations and policies Responsible stakeholders Trend: Increasing connectivity <p>Challenges for HDOs</p> <ul style="list-style-type: none"> Differences across HDO departments Effortful updating process HDO's passive position Incompatibility with existing infrastructure Long device span and legacy systems Medical device security lagging behind Overwhelming inventory management Software bundling Unexpected update behaviours Unsatisfactory vendor contact <p>Patching process steps:</p> <p>(i) Decision-making: Internal vs external</p> <ul style="list-style-type: none"> Control over maintenance Costs of inhouse maintenance Costs of outsourced maintenance Ecosystem benefits Efficiency considerations Interesting task Maintenance complexity and safety <p>(iii) Deciding to update</p> <ul style="list-style-type: none"> Avoid changes to device Cost per update/upgrade Effortful installation process Inquire with medical colleagues Inquire with technical colleagues Non-security aspects Security and safety updates always Update scope <p>(v) Installation</p> <ul style="list-style-type: none"> Installation location Installation medium Installation timing Patient care during installation Installation timeline 	<p>HDO mitigation actions</p> <ul style="list-style-type: none"> Device configuration Disconnect device Network-level mitigation Physical security measures Processes and policies Risk management <p>Risk perceptions</p> <ul style="list-style-type: none"> Hospital/patient data at risk Lack and/or delay of security updates Manageable medical device security risk Medical devices as an entry point Remotely accessible medical devices <p>Emergency updates</p> <ul style="list-style-type: none"> Prevalence Installation timeline Obligatory nature Safety priority <p>(ii) Learning about updates</p> <ul style="list-style-type: none"> Active checking on vendor platform Approach manufacturer By medical device users Direct notification by authorities Direct notification by vendor During external technician visit During sales meetings Network surveillance identifies vulnerability No active checking <p>(iv) Preparing an update</p> <ul style="list-style-type: none"> Arrangement with external technicians Arrangement with medical departments Back up plan External testing In-house testing <p>(vi) Handling post-update issues</p> <ul style="list-style-type: none"> Manufacturer post-installation support Post-update issue prevalence Rollback
---	---

Table C.2: Codebook for the interviews with Healthcare Delivery Organizations (HDO) stakeholders.

CODEBOOK: MEDICAL DEVICE MANUFACTURERS

<p>Patching process</p> <ul style="list-style-type: none"> Device OS and architecture variability Future plan: Higher patch frequency Future plan: Standardized notification channel Global regulatory mapping Notification to HDO: Call Notification to HDO: Email Notification to HDO: Letter Notification to HDO: Proprietary portal Notification to HDO: Regular service meetings Notification to HDO: Security Advisory Notification to HDO: Security white paper Patch delivery: Field Safety Notification Patch delivery: Remote Patch delivery: Service technician Patch frequency: Field Safety Notification patch Patch frequency: Regular update releases Update bundling: Considerations Update bundling: Prevalence Validation testing <p>Customer (HDO) observations</p> <ul style="list-style-type: none"> Country differences Expectations: Distrust device's security implementations Expectations: Faster patch releases Expectations: Higher patch frequency Expectations: Stricter security and privacy requirements Expectations: Variations among HDO departments Practice: Compensating network security measures Practice: High prevalence of legacy device use Practice: No full patch install coverage Update delivery considerations: Avoid functional changes Update delivery considerations: Control update process Update delivery considerations: Costs of updates Update delivery considerations: Distrust towards manufacturer Update delivery considerations: Ensure continued medical use Update delivery preferences: Install by service technicians Update delivery preferences: Low demand for remote install 	<p>Risk management</p> <ul style="list-style-type: none"> Mitigation: Device-level security measures Mitigation: Mandated Use: Device Configuration Mitigation: Mandated Use: Network security Mitigation: Mandated Use: Physical security Mitigation: Security Patch Process: Internal risk level determination Risk approach: FDA Risk criteria: Device context of use Risk criteria: Device network connections Risk criteria: Exploitability Risk criteria: Patient Safety Tool: Central security requirements framework Tool: SBOM Unrealistic risk scenarios <p>Challenges</p> <ul style="list-style-type: none"> Customers' distrust Customers' lack of awareness Dispatching technicians: Coordination Costs Evolving threat landscape Increasing vulnerabilities Internal negotiations for security Regulatory pressure: Continuous validation with faster releases Regulatory pressure: Costly validation testing Regulatory pressure: Design lock-in Remote update capability implementation Short time frame for emergency patch rollout
--	---

Table C.3: Codebook for the interviews with product security specialists at medical device manufacturers.



ACKNOWLEDGEMENTS

I am extremely grateful for all the support I have gotten during my PhD project and all the wonderful people I have had the opportunity to meet. Without them, this book would not have been written.

Firstly, I would like to thank my promotor Michel van Eeten. I am very thankful for having you as my promotor and for your continuous and inexhaustible advice and guidance. Regardless of the challenge at hand, you immediately had a helpful and actionable idea on how to approach it. I always felt this was not only stemming from your substantial experience in the field, but also from a very strong intuition under uncertainty. I learned a great deal by observing your approach and from your guidance and feedback, particularly about writing and communicating ideas. I also want to highlight how thankful I am for having a professor who cultivates such a strong social coherence in the team, which is not self-evident. By (co-)establishing recurring social events, after-work drinks, colloquia, and hosting the yearly Sinterklaas festivities at your house, you helped me and many others in the team to make friends for life. And of course, I want to thank you for all the advice on everything related to bicycles.

I would also like to wholeheartedly thank my daily supervisor and co-promotor, Simon Parkin. You have always been there to answer my questions or provide the support I needed, whether it was regarding a conceptual idea, previous literature, or writing. While other friends of mine pursuing their PhDs complained about talking to their supervisors only every other month or less, I could always count on you to respond in no time. You guided me tremendously on how to think conceptually and practically about security and the human factor by sharing your thoughts, feedback, and lessons in the Usable Security and Privacy course. You also helped me a lot with each of the papers, particularly in writing, applying Occam's razor ruthlessly, and introducing fresh ideas and perspectives on the topic at hand. I am truly grateful for having had you as my daily supervisor. And for visiting several amazing museums with you, as well as your lessons on British culture.

I also want to thank my third promotor, Carlos Gañán. You helped tremendously in writing the papers and navigating the peer-review process. Your in-depth knowledge of network security and the technical aspects of IoT security was also essential for me to plan and conduct this research. Above that, you were always on top of any administrative process required of me, usually before I was even aware of it. And thanks to your optimistic nature, I could always mentally prepare for a reject of my paper.

When starting my PhD project in the aftermath of the pandemic in 2021, I was socially deprived after several lockdowns. Additionally, I did not know anyone in the Randstad area. Thus, I am extremely grateful for meeting my fellow PhDs when coming here. Elsa, Mathew, and Arwa, you made me feel at home in no time, and you shared countless invaluable pieces of advice and wisdom with me: from how to overcome bureaucratic hurdles, mentally survive the peer-review process, and where to get the best coffee and Pita on campus. I am very grateful for the numerous lunch and coffee breaks we had together, which were as important a lesson for me in how to "do a PhD" as they could be.

I am also extremely grateful for sharing the majority of the PhD journey with my colleagues Aksel, Szu-Chun, Sandra, Swaathi, Veerle, Ronak, Evi, Kelvin, Cecile, and Fieke. Being in a similar position in our trajectory, you provided mental support, valuable ideas, and feedback on my work. Here, I would especially like to thank Aksel for helping me with a data analysis challenge during

the first study and for adding more Star Wars vibes to our office. I am also especially grateful for the experience of the Summer School at Obergurgl, Austria, in 2023, and to everyone who hosted it, as we had the opportunity to get to know each other even better.

My thanks also go to Rolf, Jury, and Savvas for their senior-level advice and guidance on many different topics: on legal and admin matters, career advice, thinking about research and academia more generally (also the art of not taking it too seriously), and for cultivating such a collegial and close atmosphere between faculty and PhDs during Thursday's drinks and beyond. And for sharing a lot of coffee capsules free of charge, of course. I would also like to thank Mattis van t' Schip for being a great sparring partner for research ideas at the intersection between law and empirical research, his thoughtful and timely contributions to Chapters 2 and 3, and for hosting me in Edinburgh and Nijmegen on several occasions. Intersect was even more fun with you on board. Many thanks also to Marie-Therese Sekwenz for reviewing my dissertation for any legal weak points or plainly incorrect statements during the final stages.

I am also grateful to the Dutch Research Council (NWO) for partially funding this research as part of the INTERSECT project (Grant nr. NWA.1160.18.301) and via project THESEUS (Grant nr. NWA.1215.18.006). Within the INTERSECT project, I would also like to thank Harold Weffers for his continuous dedication and management of the project, as well as his support in financing the surveys that were inherent to this research.

I also want to thank everyone who took the time and effort to participate in this research, either when taking the surveys, the interviews, or the more conceptual calls during research planning. In this context, I would like to thank our contacts at RDI, BSI, and EZK for their valuable input and feedback to Chapter 3. A big thank you also to the people who helped me during the challenging process of finding interview participants for Study 3, especially Wim Hafkamp and Saba Hinrichs-Krapels.

Last but certainly not least, I want to thank my family. I am extremely grateful to have my partner and wife, Keti, at my side, who believed in me from day one. We went through many significant life events during the PhD journey, and throughout this time, you never lost faith in me. I could not wish for a more supportive and kind partner at my side, and I am extremely grateful that we could share the last few years in the Netherlands, growing together as individuals and as a family. On that note, I also want to thank my daughter Mathilda, who welcomes me home every evening and teaches me what truly matters (and what does not so much).

And to my parents, Karin and Jürgen, for being there for and with me. Only due to your ongoing support and love, I was even able to reach this point.

*Lorenz Kustosch
Haarlem, November 2025*

AUTHORSHIP CONTRIBUTIONS

This dissertation is based on three papers that were created collaboratively. Hence, on this page, I describe the respective contributions of the co-authors, without whom this work would not have been possible.

In the first study (chapter 2), the co-authors Simon Parkin, Carlos H. Gañán, and Michel van Eeten provided me with invaluable guidance and feedback on the general direction of the research, the conceptualization, the methodology, and the storyline. They also supported me in writing the paper, with Simon contributing throughout the paper, but especially to the discussion section, Michel van Eeten to the introduction, and Carlos H. Gañán to the results section. Mattis van t' Schip helped us tremendously with the legal conceptualization, searched for case law of liability cases involving security or privacy incidents with IoT devices, and subsequently wrote section 2.2.2. The remainder of the writing, survey and methodology implementation, data analysis, and literature review was done by me.

For the study described in Chapter 3, my supervisors, Simon Parkin, Carlos H. Gañán, and Michel van Eeten, provided me with feedback and guidance throughout the planning and implementation of the study, as well as on the written draft of the paper. Mattis van t' Schip offered his legal expertise in conceptualizing the research, gave feedback on the written sections, and wrote section 3.2.1. The majority of the legal and academic conceptualization, survey design and implementation, data analysis, and writing was conducted by me.


In the study described in Chapter 4, Michel van Eeten provided me with contacts in the health-care sector to initiate the interview recruitment process. He, as well as Simon Parkin and Carlos Gañán, furthermore gave me valuable feedback on the paper's scope, the methodology, and the written drafts. Simon Parkin shared his substantive experience with qualitative data analysis with me, contributed to the creation and refinement of the codebook, and supported me in writing the Results and Discussion sections of the paper. Michel van Eeten was also involved in writing the introduction, supporting me in creating a concise and clear entryway into the paper. The participant recruitment, interview process, qualitative analysis, and writing of most sections were conducted by me. I would also like to thank my colleague Evi Oomens, who kindly joined one of the interviews to conduct it in Dutch, as my language skills were not yet up to par at that time.

I am extremely grateful for the opportunity to collaborate with the other co-authors. Especially with my daily supervisor, Simon, who was always available when I had a question, and for sharing his tremendous experience in conducting and writing about security research with humans. My gratitude also goes to my promoter, Michel van Eeten, who was always able to see the bigger picture of this research while also guiding on a practical, implementation level, regardless of how many PhD topics he was supervising in parallel, and to Carlos Gañán for sharing his wisdom about modeling, writing, and the intricacies of the scientific review process.



LIST OF PUBLICATIONS

- **Kustosch, L.F.**, Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2023). "Measuring Up to (Reasonable) Consumer Expectations: Providing an Empirical Basis for Holding IoT Manufacturers Legally Responsible". In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)*.
- **Kustosch, L.F.**, Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2025). "Regulating Smart Device Support Periods: User Expectations and the European Cyber Resilience Act". In *Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)*.
- **Kustosch, L.F.**, Gañán, C.H., van Eeten, M.J.G., & Parkin, S.E. (2025). "Patching Up: Stakeholder Experiences of Security Updates for Connected Medical Devices". In *Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)*.



DATASETS

Table C.4: Dataset availability

Publication	Dataset(s)
Kustosch, L.F., Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2023). "Measuring Up to (Reasonable) Consumer Expectations: Providing an Empirical Basis for Holding IoT Manufacturers Legally Responsible". In <i>Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)</i> .	Survey data can not be shared publicly due to the respective data protection measures for study subjects taken as part of the Data Management Plan.
Kustosch, L.F., Gañán, C.H., van 't Schip, M., van Eeten, M.J.G., & Parkin, S.E. (2025). "Regulating Smart Device Support Periods: User Expectations and the European Cyber Resilience Act". In <i>Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)</i> .	Aggregated and anonymized survey data can be found in this online repository, along with other supplementary material: https://doi.org/10.4121/71c038e7-e991-4dcd-9729-47dd0d9250c6
Kustosch, L.F., Gañán, C.H., van Eeten, M.J.G., & Parkin, S.E. (2025). "Patching Up: Stakeholder Experiences of Security Updates for Connected Medical Devices". In <i>Proceedings of the 34th USENIX Security Symposium (USENIX Security '25)</i> .	Raw interview data such as transcripts cannot be shared publicly due to the respective data protection measures for study subjects taken as part of the Data Management Plan.



ABOUT THE AUTHOR



Lorenz Kustosch was born in Marburg, Germany in 1992. He moved to the Netherlands and studied Psychology at the Rijksuniversiteit Groningen in 2013. After doing his minor at the Eötvös Loránd University in Budapest, he graduated in 2016 (cum laude). After a summer internship at the Fraunhofer Institute IAO in Stuttgart, he felt reinforced to pursue a post-graduate education in Human Factors Engineering and conducted his Master's in the very topic at Technische Universität (TU) Berlin. During that time, he gained insights into usability and UX in varying industry projects in his capacity as an intern and research assistant. He furthermore got the opportunity to study at Osaka University and travel Japan, where he developed a deep love for the country. He graduated from TU Berlin in 2019 with his thesis studying the digital divide for purchasing cars online in Europe. After graduating, he worked as a consultant in the automotive industry, but quickly noticed he wanted to pursue a PhD. He joined Delft University

of Technology in 2021 as a PhD researcher, where he empirically studied the role of user expectations in IoT security in the context of product conformity and liability. As part of the INTERSECT project consortium, he studied users' expectations and security-related behaviours of consumer smart devices as well as medical devices deployed at hospitals, and supervised a Master's student. He is interested in ergonomics on many levels, how security plays out in practice at organizations and people's households, and in emerging social science research methods to better understand human cognition and behaviour. In his free time, he loves to cycle, swim, and cook.



