

Measuring Up to (Reasonable) Consumer Expectations: Providing an Empirical Basis for Holding IoT Manufacturers Legally Responsible

Lorenz Kustosch
TU Delft

Carlos Gañán
TU Delft

Mattis van 't Schip
Radboud University

Michel van Eeten
TU Delft

Simon Parkin
TU Delft

Abstract

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.

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 [75] and undisclosed data collection flows [27, 78]. 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 [39]).

Home users have varying ideas on who they want to take responsibility for securing the devices before they enter the consumer market [40]. In parallel, government-level policy-makers in various countries have set standards for consumer IoT security and privacy [29, 34], 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 [15, 59, 74]. 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 [67], ask a (supposedly) 'tech-savvy' friend [68], or stop using the device altogether [15]? 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 [75].

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 [92] and thus trigger liability and product conformity regulation. While there is prior research into consumer

expectations around Internet of Things (IoT) and smart devices [55, 81, 84], 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 [40]. This does not capture what can reasonably be expected once something goes wrong with devices already in the market [40, 41, 45]. 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 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.

2.1 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 [84, 95] and privacy [10, 49], but also the security concerns they would want a solution for [21, 41, 94]. Existing work conceptualizes expectations as *normative expectations* [35, 43] – 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 [31], intentions to use devices relative to utility and data sensitivity [84], and approachability of security and privacy protection solutions [45]. Normative preferences are embodied most clearly in research on the *contextual integrity* [11] of data,

regarding individuals’ privacy preferences around the appropriateness of data flows involving IoT devices [1, 3, 6, 55, 81].

Alongside normative expectations, realistic expectations have been examined, albeit in limited scope. Zhang et al. [96] 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. [38] 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. [40] 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 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 [28]. 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 [88] 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., [8]).

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., [34, 88, 89]). Among these con-

siderations, we find expectations of ‘ordinary’ or ‘reasonable’ consumers [51]. 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., [36, 62]), and regulations highlight the importance of the expectations of a ‘public at large’ [88].

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.

2.3 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 [90] or CCPA [63], or regulations aiming at securing connected devices specifically such as the EU radio equipment directive [89]. 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 [29, 33, 34].

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 [60], network-vulnerable smart security cameras [44], and vulnerable automobile software [39]. The most common response is the release of a software patch [4, 64]. Retailers serve a role as a contact point when a purchased device has problems or must be replaced [67], and this role is being recognised in some EU countries, e.g., the Netherlands [56]. Other emerging initiatives involve IoT product ‘labels’ [30, 32, 53] 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’ [54]. We revisit how the manufacturer response informs the role of other stakeholders (such as retailers and governments) in section 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 [15, 20]. Otherwise, if a user has problems with a device, they may reach out to someone they regard as ‘informal’ technical support [68], or seek information on news or specialist websites [72] or on forums [42].

3 Method

To address our research questions defined in section 1, we deployed a vignette-driven online survey with 862 participants from the Prolific [70] 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 [3, 6, 9, 12, 14, 50, 55] and allow to study participants’ judgments on multidimensional phenomena, while reducing social desirability biases of direct survey questions [46, 87]. Vignettes have been demonstrated elsewhere to be useful in exploring scenarios with *unresolved issues* in security and privacy [14], 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 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.

Factor	Levels security vignettes	Levels privacy vignettes
Device	1. Smart speaker 2. Smart watch 3. Smart washing machine 4. Smart security camera 5. Smartphone 6. Connected car	1. Smart speaker 2. Smart watch 3. Smart washing machine 4. Smart security camera 5. Smartphone 6. Connected car
Event	1. DDoS 2. Unauthorized data access 3. Ransomware	1. Data collection without consent 2. Third party data sharing 3. Forced data collection
Manufacturer response	1. Announce patch 2. Inform users via website and social media 3. No response 4. Recall	1. Announce update with more privacy settings 2. Inform users via updated privacy policy 3. No response -
User response	1. Attempt to return device 2. Attempt technical mitigation 3. Seek advice online 4. Turn device off 5. Keep using as before	1. Attempt to return device 2. Attempt technical mitigation 3. Seek advice online 4. Turn device off 5. Keep using as before

Table 1: Overview of vignette factors and levels.

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 [40];
- 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).

3.2 Survey procedure

Participants on the Prolific platform were directed to a Qualtrics [71] 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 [46]). 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 1, constructed to resemble a scenario as in the example vignette (subsection 3.1). Participants then answered questions about a differing set of these kinds of vignettes. Participants either received a full set of security vignettes, or of privacy vignettes. Vignette construction is detailed in subsection 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 [73] 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 online as accompanying material¹.

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 [7, 77]. We opted to present seven vignettes to strike a balance between more repeated measures per participant (increasing statistical power) [7] while not mentally overloading them with too many vignettes [48]. After each vignette the four Likert scale questions described in subsection 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 1 led to a total vignette population of

¹<https://doi.org/10.4121/c.6440264.v1>

360 different combinations for security, and 270 different combinations for privacy. Adhering to methodological literature [7, 82], 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 [48].

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.

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 [52], unauthorized access to IoT sensor data [23, 76], and ransomware attacks targeting IoT devices [66]. Privacy-related events included reports of staff listening to device recordings for training of algorithms [27, 57], or device data being shared without the user’s consent [26, 78].

Events followed one of three different outcomes: that continued use of a device is impaired 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 [4, 64] and privacy issues such as providing more privacy controls [17]. 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 [60], smart security cameras being vulnerable to DDoS malware [44], or smart vehicle vulnerabilities [39]. There were also accounts of manufacturers not visibly responding directly to an event [22, 61, 65], 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 [18, 19, 69], 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 [15, 74], as existing work mostly focuses on preventative mitigation by users [2, 37, 40, 41, 83]. We included five different user responses: 1) Keep using the device, due to e.g., discounting of risks to data [45, 47, 83, 94] or security [40, 83], or resignation [40, 47, 79]; 2) Unplug the device, ceasing or pausing use [15, 76, 86]; 3) Opportunistically seek help from others [24, 68] or online [5, 42, 76, 80]; 4) Attempt technical remediation oneself through device configuration or isolation from the network [37, 41]; 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.

3.4.1 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 [15, 76, 86].

3.4.2 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 [70] 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.

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.

3.6 Participants

Participants were recruited via the crowdsourcing platform Prolific [70] 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. 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 [85]. 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

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

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.

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 [13]. 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 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).

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 [7], 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 [58].

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 3 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 [16], initial codes of reoccurring themes in the data were generated, which were then regularly discussed between the researchers in an iterative coding process.

4 Results

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.

4.1.1 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 1 shows how likely the manufacturer responses to a security event were rated on average across device types; Figure 2 does the same for privacy events.

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 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 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 3, Model 3) supported this trend: for security, all manufacturer responses were judged as significantly more likely than no response.

Figure 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 3, Model 4) shows that the likelihoods of the two explicit manufacturer responses were comparable, indicated by similar coefficient estimates.

4.1.2 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 [40]. The right-hand side of both Figure 1 and Figure 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 3 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 manufacturer 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 3 (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), releasing a software update with more privacy controls was most preferred across devices and privacy events (mean of block '*Announce update with*

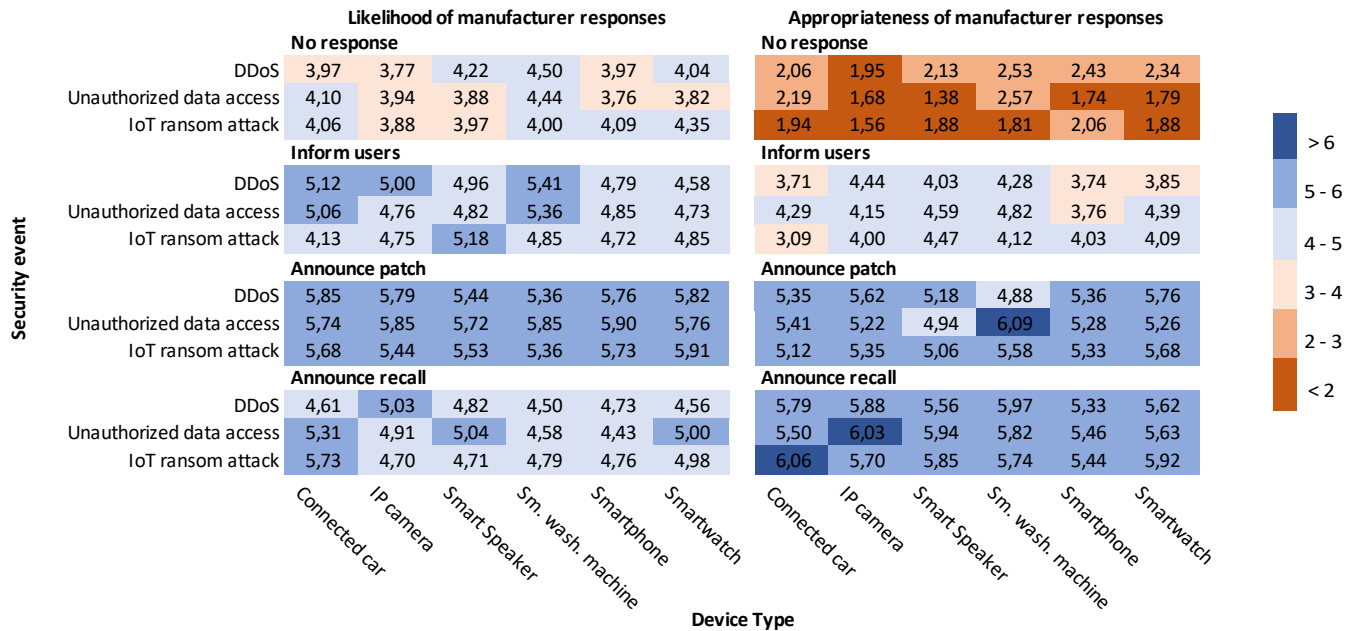


Figure 1: Ratings of likelihood and appropriateness of manufacturer responses to security events. Measured on a 7-point Likert scale with 1 = ‘Extremely unlikely’ to 7 = ‘Extremely likely’ for likelihood of manufacturer response and 1 = ‘Strongly disagree’ to 7 = ‘Strongly agree’ for appropriateness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

privacy settings’ = 5.09). The regression model predicting the appropriateness ratings for privacy (Table 3, Model 2) supported this, as both of the explicit manufacturer responses were rated significantly more appropriate than no response.

4.1.3 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 1 and Figure 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).

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. Regression Model 3 in Table 3 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 smartphones 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) negatively impacted how appropriate manufacturer responses were rated, especially for *‘No response’*, as also shown in the regression model (Table 3, Model 2, sign. diff. between no consent and third-party sharing).

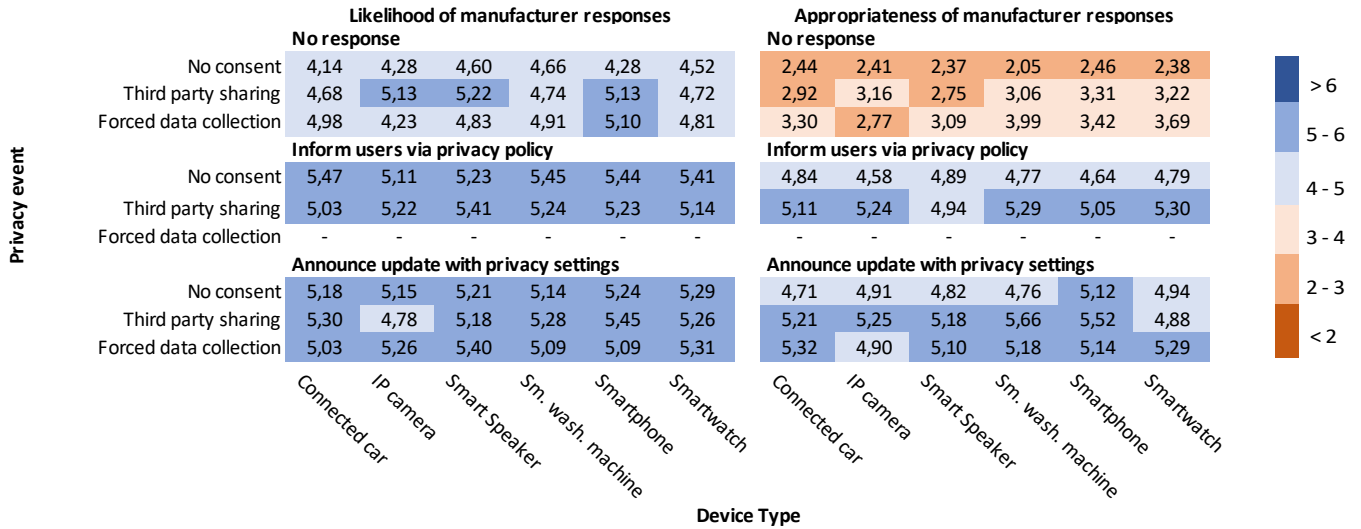


Figure 2: Ratings of likelihood and appropriateness of manufacturer responses to privacy events. Measured on a 7-point Likert scale with 1 = ‘Extremely unlikely’ to 7 = ‘Extremely likely’ for likelihood of manufacturer response and 1 = ‘Strongly disagree’ to 7 = ‘Strongly agree’ for appropriateness of manufacturer response. Empty cells correspond to less plausible vignette level combinations, which were removed from the design.

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 judgments translate to their expectations about the suitability of the user’s options in reaching a satisfactory response to particular events. **Figure 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 3** present the results of the regression models predicting the suitability of the user responses with vignette factors.

4.2.1 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 refund’ = 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 3** (**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 attempt-

ing 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 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 3**, **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).

4.2.2 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.

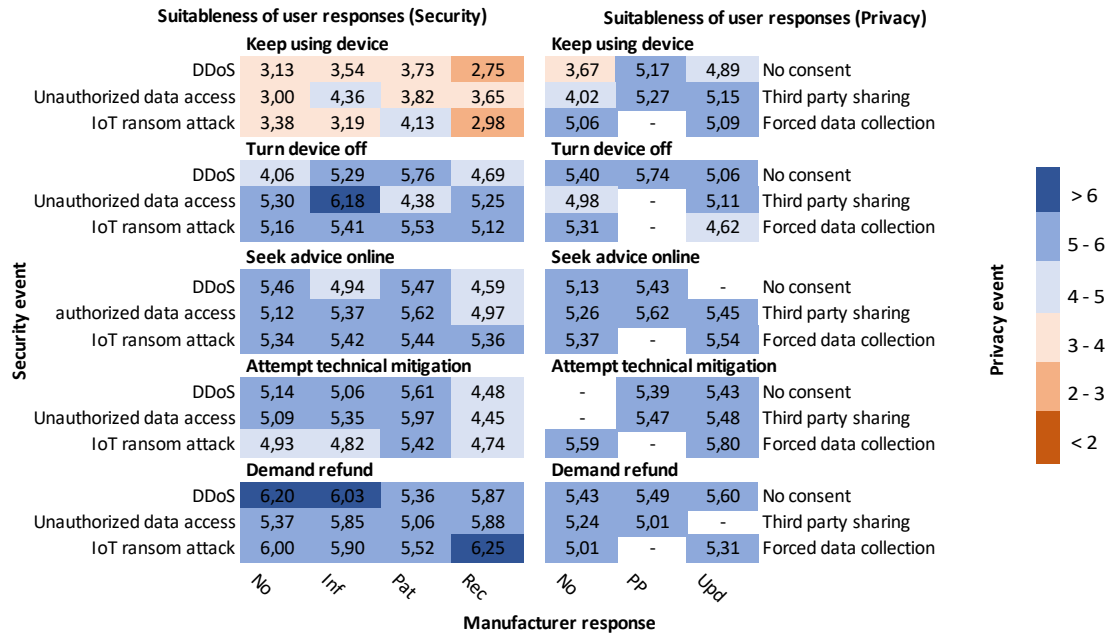


Figure 3: Ratings of user response suitability across security and privacy events and previous manufacturer responses. Measured on a 7-point Likert scale with 1 = ‘Strongly disagree’ and 7 = ‘Strongly agree’. Empty cells correspond to less plausible combinations, which were removed from the design or 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’.

There were lower regression coefficient estimates for the user responses to privacy events (*Model 6*) than for security events (*Model 5*) (see [Table 3](#)): 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.

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 re-setting. 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 [15, 47, 74].

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 4.2.2, as they demonstrate a general uncertainty about data flows and how IoT users should manage privacy.

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 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 [34], and highlights the position of ISPs as being able to triangulate security problems to specific users (e.g., [15]).

The picture was less clear for emerging privacy issues with IoT devices (Figure 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 [63, 90] and a learned helplessness and resignation regarding control over the occurrence of privacy violations, and is in accordance with prior work [40, 47, 79].

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, 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., [92]) 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 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 [15, 74]. 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 [91]).

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) (reinforcing prior findings elsewhere [93]). 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 [93], 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 [40]. 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., [40, 41, 55, 84]) 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 [32, 84]. 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 [3, 6, 21, 55, 80], 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.

How best to involve users. For RQ4, the user's involvement in addressing security and privacy risks was assessed (Figure 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' [93] 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 [15] and privacy [41, 45]), 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 [15, 76, 86], 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 subsection 4.2.2), 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 protect-

ing 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 [40]. 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.

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 [34] 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 [40], 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.

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 [85]. 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 [32, 72], such as word-of-mouth, unusual device behaviour, or direct notifications (e.g., by ISPs [20]). We opted for the news post as this is a common channel for home users [25, 72] and may be communicated itself by word-of-mouth or analogy [73]. 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 4.3](#).

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.

Acknowledgments

This work has been partially supported by the INTERSCT project, Grant No. NWA.1160.18.301, funded by Netherlands Organisation for Scientific Research (NWO). The findings reported herein are solely responsibility of the authors.

References

- [1] Denielle Abaquita, Paritosh Bahirat, Karla A Badillo-Urquiola, and Pamela Wisniewski. Privacy norms within the internet of things using contextual integrity. In *Companion of the 2020 ACM International Conference on Supporting Group Work*, pages 131–134, 2020.
- [2] Noura Abdi, Kopo M Ramokapane, and Jose M Such. More than smart speakers: security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 451–466, 2019.
- [3] Noura Abdi, Xiao Zhan, Kopo M Ramokapane, and Jose Such. Privacy norms for smart home personal assistants. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–14, 2021.
- [4] Lawrence Abrams. QNAP force-installs update after DeadBolt ransomware hits 3,600 devices, 2022. URL: <https://www.bleepingcomputer.com/news/security/qnap-force-installs-update-after-deadbolt-ransomware-hits-3-600-devices/>.
- [5] Sara Amini and Chris Kanich. Characterizing the impact of malware infections and remediation attempts through support forum analysis. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*, pages 70–78, 2017. doi:10.1109/ECRIME.2017.7945056.
- [6] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–23, 2018. doi:10.1145/3214262.
- [7] Katrin Auspurg and Thomas Hinz. *Factorial survey experiments*, volume 175. Sage Publications, 2014.
- [8] An Baeyens and Tom Goffin. Boston Scientific Medizintechnik GmbH v. AOK Sachsen-Anhalt. *European journal of health law*, 22(3):301–307, 2015. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62013CJ0503>.
- [9] Paritosh Bahirat, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. 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*, pages 1–18. ACM, 2021. doi:10.1145/3411764.3445672.
- [10] Natã M Barbosa, Zhuohao Zhang, and Yang Wang. 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)*, pages 417–435, 2020.
- [11] A. Barth, A. Datta, J.C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: framework and applications. In *2006 IEEE Symposium on Security and Privacy (S&P'06)*, pages 15 pp.–198, 2006. doi:10.1109/SP.2006.32.
- [12] Adam Beutement, Ingolf Becker, Simon Parkin, Kat Krol, and Angela Sasse. Productive security: A scalable methodology for analysing employee security behaviours. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 253–270, 2016.
- [13] Ingolf Becker, Simon Parkin, and M Angela Sasse. Measuring the Success of Context-Aware Security Behaviour Surveys. In *The LASER Workshop: Learning from Authoritative Security Experiment Results (LASER 2017)*, pages 77–86, 2017.
- [14] John M. Blythe, Lynne Coventry, and Linda Little. 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*, pages 103–122. USENIX Association, 2015.
- [15] Brennen Bouwmeester, Elsa Rodríguez, Carlos Gañán, Michel van Eeten, and Simon Parkin. “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)*, pages 493–512, 2021.
- [16] Virginia Braun and Victoria Clarke. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology*, 18(3):328–352, 2021.
- [17] Michael Brown. Ring announces 6 new products, along with measures to quiet criticism of its privacy practices, 2020. URL: <https://www.techhive.com/article/584254/ring-announces-6-new-products-along-with-measures-to-quiet-criticism-of-its-privacy-practices.html>.
- [18] Jonathan Bundy and Michael D. Pfarrer. A Burden of Responsibility: The Role of Social Approval at the Onset of a Crisis. *Academy of Management Review*, 40(3):345–369, 2015. doi:10.5465/amr.2013.0027.
- [19] Jonathan Bundy, Michael D. Pfarrer, Cole E. Short, and W. Timothy Coombs. Crises and Crisis Management: Integration, Interpretation, and Research Development. *Journal of Management*, 43(6):1661–1692, 2017. doi:10.1177/0149206316680030.
- [20] Orçun Çetin, Carlos Ganán, Lisette Altena, Takahiro Kasama, Daisuke Inoue, Kazuki Tamiya, Ying Tie, Katsunari Yoshioka, and Michel Van Eeten. Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. In *NDSS*, 2019.
- [21] George Chalhouh, Martin J Kraemer, Norbert Nthala, and Ivan Flechais. “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*, pages 1–16. ACM, 2021. doi:10.1145/3411764.3445691.
- [22] Catalin Cimpanu. Hackers are hijacking smart building access systems to launch DDoS attacks, 2020. URL: <https://www.zdnet.com/article/hackers-are-hijacking-smart-building-access-systems-to-launch-ddos-attacks/>.
- [23] Ry Crist. ADT technician pleads guilty to spying on customer camera feeds for years, 2021. URL: <https://www.cnet.com/home/smart-home/adt-home-security-technician-pleads-guilty-to-spying-on-customer-camera-feeds-for-years/>.

- [24] Sauvik Das, Laura A Dabbish, and Jason I Hong. A Typology of Perceived Triggers for End-User Security and Privacy Behaviors. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 97–115, 2019.
- [25] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. 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*, pages 1–12, 2018.
- [26] Wendy Davis. Vizio Hit With Privacy Lawsuit Over Connected TVs, 2015. URL: <https://www.mediapost.com/publications/article/262835/vizio-hit-with-privacy-lawsuit-over-connected-tvs.html>.
- [27] Mattt Day, Giles Turner, and Natalia Drozdiak. Thousands of Amazon Workers Listen to Alexa Users’ Conversations, 2019. URL: <https://time.com/5568815/amazon-workers-listen-to-alexa/>.
- [28] Benjamin Dean. An Exploration of Strict Products Liability and the Internet of Things. Technical report, Center for Democracy and Technology, 2018. URL: <https://cdt.org/wp-content/uploads/2018/04/2018-04-16-IoT-Strict-Products-Liability-FNL.pdf>.
- [29] Department for Digital, Culture, Media and Sport (UK Government). Code of Practice for Consumer IoT Security. 2018. URL: <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security>.
- [30] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. Ask the Experts: What Should Be on an IoT Privacy and Security Label? In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 447–464, 2020. doi:10.1109/SP40000.2020.00043.
- [31] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 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)*, pages 519–536, 2021. doi:10.1109/SP40001.2021.00112.
- [32] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12. ACM, 2019. doi:10.1145/3290605.3300764.
- [33] ENISA. Good Practices for Security of IoT - Secure Software Development Lifecycle. 2019. URL: <https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot-1>.
- [34] Michael Fagan, Katerina Megas, Karen Scarfone, and Matthew Smith. Foundational Cybersecurity Activities for IoT Device Manufacturers. Technical Report NIST Internal or Interagency Report (NISTIR) 8259, National Institute of Standards and Technology, 2020. URL: <https://csrc.nist.gov/publications/detail/nistir/8259/final>, doi:10.6028/NIST.IR.8259.
- [35] Nathan Favero and Minjung Kim. Everything Is Relative: How Citizens Form and Use Expectations in Evaluating Services. *Journal of Public Administration Research and Theory*, 31(3):561–577, 2021. doi:10.1093/jopart/muaa048.
- [36] District Court for the Western District of Pennsylvania. Cohen v. Johnson & Johnson, 2022. URL: <https://law.justia.com/cases/federal/district-courts/pennsylvania/pawdce/2:2020cv00057/262978/68/>.
- [37] Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman. Privacy and security threat models and mitigation strategies of older adults. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 21–40, 2019.
- [38] Sandra Gabriele and Sonia Chiasson. Understanding Fitness Tracker Users’ Security and Privacy Knowledge, Attitudes and Behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12. ACM, 2020. doi:10.1145/3313831.3376651.
- [39] Andy Greenberg. After Jeep Hack, Chrysler Recalls 1.4M Vehicles for Bug Fix. *Wired*, 2015. URL: <https://www.wired.com/2015/07/jeep-hack-chrysler-recalls-1-4m-vehicles-bug-fix/>.
- [40] Julie Haney, Yasemin Acar, and Susanne Furman. “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)*, pages 411–428, 2021.
- [41] Julie M. Haney, Susanne M. Furman, and Yasemin Acar. Smart home security and privacy mitigations: Consumer perceptions, practices, and challenges. In Abbas Moallem, editor, *HCI for Cybersecurity, Privacy and Trust*, Lecture Notes in Computer Science, pages 393–411. Springer International Publishing, 2020. doi:10.1007/978-3-030-50309-3_26.
- [42] Ayako A Hasegawa, Naomi Yamashita, Tatsuya Mori, Daisuke Inoue, and Mitsuaki Akiyama. Understanding Non-Experts’ Security- and Privacy-Related Questions on a Q&A Site. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 39–56, 2022.
- [43] Bronwyn Higgs, Michael Jay Polonsky, and Mary Hollick. Measuring expectations: forecast vs. ideal expectations. does it really matter? *Journal of retailing and consumer services*, 12(1):49–64, 2005.
- [44] Sijia Jiang and Jim Finkle. China’s Xiongmai to recall up to 10,000 webcams after hack. *Reuters*, 2016. URL: <https://www.reuters.com/article/us-cyber-attacks-china-idUSKCN12P1TT>.
- [45] Haojian Jin, Boyuan Guo, Rituparna Roychoudhury, Yaxing Yao, Swarun Kumar, Yuvraj Agarwal, and Jason I. Hong. Exploring the Needs of Users for Supporting Privacy-Protective Behaviors in Smart Homes. In *CHI Conference on Human Factors in Computing Systems*, pages 1–19. ACM, 2022. doi:10.1145/3491102.3517602.
- [46] Kat Krol, Jonathan M Spring, Simon Parkin, and M Angela Sasse. Towards robust experimental design for user studies in security and privacy. In *The LASER Workshop: Learning from Authoritative Security Experiment Results (LASER 2016)*, pages 21–31, 2016.

- [47] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–31, 2018. doi:10.1145/3274371.
- [48] Ye Li, Antonia Krefeld-Schwalb, Daniel G. Wall, Eric J. Johnson, Olivier Toubia, and Daniel M. Bartels. The More You Ask, the Less You Get: When Additional Questions Hurt External Validity. *Journal of Marketing Research*, 2021. doi:10.1177/00222437211073581.
- [49] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. "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*, pages 1–11, 2020.
- [50] Kirsten Martin and Katie Shilton. 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):1871–1882, 2016. doi:10.1002/asi.23500.
- [51] Clayton J Masterman and W Kip Viscusi. The specific consumer expectations test for product defects. *Ind. LJ*, 95:183, 2020.
- [52] Joseph Menn, Jim Finkle, and Dustin Volz. Cyber attacks disrupt PayPal, Twitter, other sites. *Reuters*, 2016. URL: <https://www.reuters.com/article/us-usa-cyber-idUSKCN12L1ME>.
- [53] Philipp Morgner, Christoph Mai, Nicole Koschate-Fischer, Felix Freiling, and Zinaida Benenson. Security Update Labels: Establishing Economic Incentives for Security Patching of IoT Consumer Products. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 429–446, 2020. doi:10.1109/SP40000.2020.00021.
- [54] Mozilla. *Privacy not included. <https://foundation.mozilla.org/en/privacynotincluded/>, 2021.
- [55] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. Privacy expectations and preferences in an IoT world. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 399–412, 2017.
- [56] RVO Netherlands Enterprise Agency. Supplying updates is mandatory. 2022. URL: <https://business.gov.nl/amendment/supplying-updates-to-become-mandatory/>.
- [57] BBC News. Google seeks permission for staff to listen to Assistant recordings. *BBC News*, 2019. URL: <https://www.bbc.com/news/technology-49796207>.
- [58] Geoff Norman. Likert scales, levels of measurement and the "laws" of statistics. *Advances in Health Sciences Education*, 15(5):625–632, 2010. doi:10.1007/s10459-010-9222-y.
- [59] Norbert Nthala and Ivan Flechais. Informal support networks: an investigation into home data security practices. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pages 63–82, 2018.
- [60] Lindsey O'Donnell. EU Recalls Children's Smartwatch That Leaks Location Data, 2019. URL: <https://threatpost.com/eu-recalls-childrens-smartwatch-that-leaks-location-data/141511/>.
- [61] Lindsey O'Donnell. Unpatched Security Flaws Open Connected Vacuum to Takeover, 2020. URL: <https://threatpost.com/unpatched-security-flaws-open-connected-vacuum-to-takeover/153142/>.
- [62] District Court of Appeal of the State of Florida. Grieco v. Daiho Sangyo, Inc., 2022. URL: <https://law.justia.com/cases/florida/fourth-district-court-of-appeal/2022/20-2294.html>.
- [63] State of California. California Consumer Privacy Act. 2018. URL: https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5.
- [64] Charlie Osborne. Remote code execution flaw allowed hijack of Motorola Halo+ baby monitors, 2021. URL: <https://portswigger.net/daily-swig/remote-code-execution-flaw-allowed-hijack-of-motorola-halo-baby-monitors>.
- [65] Charlie Osborne. Researcher discloses alleged zero-day vulnerabilities in NUUO NVRmini2 recording device, 2022. URL: <https://portswigger.net/daily-swig/researcher-discloses-alleged-zero-day-vulnerabilities-in-nuuo-nvrmini2-recording-device>.
- [66] Pierluigi Paganini. Watch out, FLocker Ransomware targets Android smart TVs, 2016. URL: <https://securityaffairs.co/wordpress/48383/iot/flocker-ransomware-smart-tvs.html>.
- [67] Simon Parkin, Elissa M Redmiles, Lynne Coventry, and M Angela Sasse. 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, 2019.
- [68] Erika Shehan Poole, Marshini Chetty, Tom Morgan, Rebecca E Grinter, and W Keith Edwards. Computer help at home: methods and motivations for informal technical support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 739–748, 2009.
- [69] Lutz Preuss, Ralf Barkemeyer, Olivier Gergaud, and Christophe Faugère. Corporate Scandals: Firm Response Strategies and Subsequent Media Coverage. *Academy of Management Proceedings*, 2018(1):18369, 2018. doi:10.5465/AMBPP.2018.18369abstract.
- [70] Prolific. URL: <https://www.prolific.co/>.
- [71] Qualtrics. URL: <https://www.qualtrics.com>.
- [72] Emilee Rader and Rick Wash. Identifying patterns in informal sources of security information. *Journal of Cybersecurity*, 1(1):121–144, 2015.
- [73] Emilee Rader, Rick Wash, and Brandon Brooks. Stories as informal lessons about security. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, pages 1–17, 2012.

- [74] Elsa Rodríguez, Max Fukkink, Simon Parkin, Michel van Eeten, and Carlos Gañán. Difficult for thee, but not for me: Measuring the difficulty and user experience of remediating persistent iot malware. *arXiv preprint arXiv:2203.01683*, 2022.
- [75] Elsa Rodriguez, Arman Noroozian, Michel van Eeten, and Carlos Gañán. Super-spreaders: Quantifying the role of iot manufacturers in device infections. In *20th Annual Workshop on the Economics of Information Security (WEIS 2021)*, 2021.
- [76] Asreen Rostami, Minna Vigren, Shahid Raza, and Barry Brown. Being hacked: Understanding victims' experiences of iot hacking. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 613–631, 2022.
- [77] Carsten Sauer, Katrin Auspurg, and Thomas Hinz. 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):195–214, 2020.
- [78] Bruce Schneier. Ring Gives Videos to Police without a Warrant or User Consent - Schneier on Security, 2022. URL: <https://www.schneier.com/blog/archives/2022/08/ring-gives-videos-to-police-without-a-warrant-or-user-consent.html>.
- [79] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 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*, pages 2347–2356. ACM, 2014. doi: 10.1145/2556288.2557421.
- [80] Yan Shvartzshnaider, Madelyn Rose Sanfilippo, and Noah Aphorpe. Contextual Integrity as a Gauge for Governing Knowledge Commons. In Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg, editors, *Governing Privacy in Knowledge Commons*, Cambridge Studies on Governing Knowledge Commons, pages 220–244. Cambridge University Press, 2021. doi:10.1017/9781108749978.010.
- [81] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. Learning privacy expectations by crowdsourcing contextual informational norms. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 4:209–218, 2016.
- [82] Dan Su and Peter M. Steiner. An Evaluation of Experimental Designs for Constructing Vignette Sets in Factorial Surveys. *Sociological Methods & Research*, 49(2):455–497, 2020. doi: 10.1177/0049124117746427.
- [83] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. “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)*, pages 435–450, 2019.
- [84] Madiha Tabassum, Tomasz Kosiński, Alisa Frik, Nathan Malkin, Primal Wijesekera, Serge Egelman, and Heather Richter Lipford. Investigating Users’ Preferences and Expectations for Always-Listening Voice Assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(4):1–23, 2019. doi:10.1145/3369807.
- [85] Jenny Tang, Eleanor Birrell, and Ada Lerner. 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)*, pages 367–385, 2022.
- [86] Huixin Tian, Chris Kanich, Jason Polakis, and Sameer Patil. Tech Pains: Characterizations of Lived Cybersecurity Experiences. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 250–259, 2020. doi:10.1109/EuroSPW51379.2020.00040.
- [87] Alan J. Tomassetti, Reeshad S. Dalal, and Seth A. Kaplan. Is Policy Capturing Really More Resistant Than Traditional Self-Report Techniques to Socially Desirable Responding? *Organizational Research Methods*, 19(2):255–285, 2016. doi: 10.1177/1094428115627497.
- [88] European Union. Product Liability Directive. 1985. URL: <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31985L0374:en:HTML>.
- [89] European Union. Radio Equipment Directive. 2014. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32014L0053>.
- [90] European Union. General Data Protection Regulation. 2016. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
- [91] European Union. Questions and Answers: Strengthening cybersecurity of wireless devices and products, 2021. URL: https://ec.europa.eu/commission/presscorner/detail/en/QANDA_21_5635.
- [92] Daniel W Woods and Aaron Ceros. Blessed are the lawyers, for they shall inherit cybersecurity. In *New Security Paradigms Workshop*, pages 1–12, 2021.
- [93] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 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*, pages 1–12, 2019.
- [94] Eric Zeng, Shrirang Mare, and Franziska Roesner. End user security and privacy concerns with smart homes. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 65–80, 2017.
- [95] Eric Zeng and Franziska Roesner. 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)*, pages 159–176. USENIX Association, 2019.
- [96] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh. “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), 2021. doi:10.2478/popets-2021-0028.

B Regression tables	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
Coefficient	Estimate	Estimate	Estimate	Estimate	Estimate	Estimate
Intercept	1.65 ***	2.61 ***	3.80 ***	4.52 ***	3.15 ***	4.63 ***
Device:						
Connected car	0.30	-0.03	0.27	0.05	-0.35	-0.34***
Smart speaker	0.05	-0.07	0.19	0.20 *	0.06	-0.12
Smart washing machine	0.64 ***	-0.08	0.54 ***	0.12	-0.26	-0.03
Smartphone	0.39 *	-0.07	0.15 *	0.21 **	0.07	-0.22**
Smartwatch	0.32	-0.04	0.28	0.13	0.36 *	-0.06
Security camera (reference)						
Security/ Privacy Event:						
IoT ransomware	-0.10		-0.02		0.16 **	
Unauthorized data access	-0.01		0.05		0.08	
DDoS (reference)						
Forced data collection		0.19		0.16 *		0.47 ***
Third party sharing		0.44 **		0.15 **		-0.04
No consent (reference)						
Manufacturer response:						
(Sec) Inform users	2.37 ***		1.02 ***		0.30	
(Sec) Recall	4.21 ***		1.14 ***		-0.17	
(Sec) Patch	3.54 ***		1.89 ***		0.45 *	
(Sec) No response (reference)						
(Priv) Inform via privacy policy		2.14***		0.56 ***	2.01 ***	0.34 ***
(Priv) Privacy S-W update		2.12 ***		0.47 ***	2.65 ***	0.58 ***
(Priv) No response (reference)					0.87 ***	0.13
User response:					1.73 ***	0.12
Advice online forums					1.95 ***	0.77 ***
Return device						0.07
Tech. mitigation						0.18 **
Turn off						0.06
Keep using (reference)						
Interaction Effects:						
ManuResp*Device:		Device*PrivEvent:	ManuResp*Device:	None found	ManuResp*Device:	Event*UserResp:
Inform*Car (-0.75**)		Car*Forced (0.43**)	Inform*Smartwatch (-0.81**)		Inform*Smartwatch (-0.81**)	Forced*Return (-0.82***)
WashMachi (-0.51*)		WashMachi*Forced (0.78*)	Recall*Car (0.54*)		Recall*Car (0.54*)	Forced*TurnOff (-0.98***)
Recall*WashMachi(-0.65**)		Smartphone*Forced (0.57**)	Patch*WashMachi (-0.69*)		Patch*Smartphone (-0.98***)	
Smartphone (-0.63)		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 3: 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’ and 7 = ‘Strongly agree’ to the statement ‘The manufacturer’s response to the 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.