



Examining Consumer Reviews to Understand Security and Privacy Issues in the Market of Smart Home Devices

Swaathi Vetrivel, Veerle van Harten, Carlos H. Gañán, Michel van Eeten,
and Simon Parkin, *Delft University of Technology*

<https://www.usenix.org/conference/usenixsecurity23/presentation/vetrivel>

This paper is included in the Proceedings of the
32nd USENIX Security Symposium.

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.

Examining Consumer Reviews to Understand Security and Privacy Issues in the Market of Smart Home Devices

Swaathi Vetrivel, Veerle van Harten, Carlos H. Gañán, Michel van Eeten, and Simon Parkin

Delft University of Technology

{S.Vetrivel, V.T.C.vanHarten, C.HernandezGanan, M.J.G.vanEeten,
S.E.Parkin}@tudelft.nl

Abstract

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

1 Introduction

Among the range of consumer IoT devices now available - like smart doorbells and smart home surveillance systems - many lack sufficient security and privacy features that are fit for purpose. These shortcomings have been exploited in various ways, most visibly in large-scale Distributed Denial of Service (DDoS) attacks from compromised IoT devices [1,2]. These emerge as part of a broader trend to leverage vulnerable IoT devices for malicious activities, ranging from botnets as criminal infrastructure [3] to cryptojacking [4], and intimate domestic abuse [5,6].

A broad consensus has emerged that the root cause of this trend is a market failure caused by incentive misalignment [7].

Manufacturers of the affected devices do not have sufficient incentive to improve security. The cost of poor security is not borne by them but by others, such as the network operators and service providers that suffer DDoS attacks. However, the incentives for manufacturers would improve if customers care about and prioritise security and privacy in their IoT device purchases - and recent research demonstrates that users not only care about IoT security and privacy but are also willing to pay for it [8,9]. The failure then is one of information asymmetry, a 'market for lemons' [10] where the consumer cannot discern good from bad with the information available.

In order to fix this market failure and hasten improvements to the security and privacy ('S&P' from here onwards) of IoT devices, external interventions have been pursued. To that end, there are various efforts to introduce a range of *market signals* [11,12] - more information about the quality of the device - to reduce the *information asymmetry* between S&P observers (such as experts, governments) and consumers. These signals can be marketing-controlled, within the control of the manufacturer like security 'labels', or non-marketing-controlled like customer reviews [13].

Self-certified security labels that state assurances [14-16] and aim to standardise the S&P information available to consumers - like food nutrition labelling - are still in formative stages and are not yet in use. Further, it is not assured that they would reflect the actual real-world security posture of the device or directly respond to the concerns that consumers have [17]. Crucially, these efforts overlook how much the consumer base is already recognising - and signalling - a need for S&P in IoT devices and what the expression of those needs looks like. Such signals are present within non-marketing-controlled avenues like customer reviews.

Our primary motivation in this study is to evaluate if customer reviews - as an existing source of signals about the quality of consumer IoT products - voice consumers' S&P-related concerns. Given that prospective buyers look at reviews when considering a purchase [18], understanding the content of reviews can inform the views that consumers express and also how accessible those views are to other consumers as signals

of device quality. In turn, this can point to where action is needed to further inform decisions around the purchase of a smart device.

Moreover, if S&P issues raised in customer reviews can be characterised and amplified, they will highlight the S&P posture of IoT devices and potentially incentivise manufacturers through the impact on brand reputation [19]. These signalling effects are agnostic to whether the S&P-issues in the reviews are trustworthy, misinformed, or even whether the review is fake [20], as long as prospective buyers trust the reviews as an important source of information which research suggests they do [18].

Here, we present a large-scale investigation and characterisation of signals for S&P, through an analysis of consumer reviews, in the market of IoT devices commonly infected with ‘Mirai-like’ malware [1, 21] and related device features. This device selection allowed us to see to what extent S&P-issues notorious within the security community have permeated the marketplace and are raised – unprompted – by consumers.

We address two research questions: **(RQ1)** What fraction of customer reviews for IoT devices articulate security or privacy issues? **(RQ2)** When security issues or privacy are mentioned, what themes are being articulated? To answer these questions, via distinct sampling strategies, topic modelling and qualitative analysis approaches, we collected and analysed 83,686 reviews from all of the six country websites of Amazon, which are natively in English. The scraping process from product link collection, cleanup and subsequent review data collection was conducted between May and August 2021.

Our main contributions are as follows:

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

- We discuss several recommendations for strengthening the information and signalling value of customer reviews and leveraging this existing mechanism to reduce information asymmetry and the security incentives of manufacturers.

2 Background and Related Work

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

2.1 Analysis of online reviews

Online product reviews have been recognised as containing critical information regarding consumers’ concerns [22] and are considered essential in building a firm’s business intelligence [23]. Reviews have been noted to function both as informants and as recommenders influencing both product sales and purchase decisions [24]. Moreover, studies on the impact of online reviews show that the quality and quantity of product reviews positively influence the sales of a product [25]. With respect to reviewers themselves, Hu et al. [26] show the self-selection bias at play where buyers with extreme positive or negative experiences are more likely to post online reviews. However, Han et al. [27] report that this bias is mitigated when buyers are familiar with the online review platform. Given the popularity of Amazon, we expect that most buyers are familiar with its review platform and are, therefore, more likely to post reviews for varying levels of satisfaction. This could also explain why Amazon is the most popular source of customer reviews within the research community [28].

Moreover, customer reviews of other devices, including certain types of IoT devices like smart home assistants and wearables, have been studied previously. Two such studies [29, 30] use unsupervised machine learning techniques to understand if users express any privacy concerns in reviews of popular e-commerce sites, including Amazon. The results vary between the studies from a significant percentage of users concerned with privacy [29] to only 2% [30]. Using thematic analysis, Linden et al. [31] performed a comparative analysis of customer reviews, also on Amazon, of human and pet wearables and found that very few privacy concerns were expressed about these technologies. However, in our work, we go beyond merely reporting S&P concerns, quantifying the presence of S&P issues in reviews and also qualitatively determining the character of S&P signals within these reviews.

2.2 Consumer perceptions of IoT security and privacy

Earlier work on consumer perceptions of IoT S&P has primarily been conducted through surveys, semi-structured in-

interviews and experience sampling. To examine the mental models of users of smart devices, Abdi et al. [32] and Zeng et al. [33] conducted semi-structured interviews and reported gaps in users' mental models regarding security. They attribute these gaps primarily to limited technical understanding and point to ad-hoc (and typically non-technical) strategies employed by users in order to protect themselves.

With respect to privacy, a study by Williams et al. [34] observes that the price of consumer IoT devices deterred more users than privacy concerns and note that since the purchase of IoT devices is voluntary, privacy was more likely to be sacrificed for functionality rather than necessity. This trade-off is echoed in other studies [35, 36], which observe that users balance the risks of using IoT devices against the convenience and benefits offered. In our analysis of reviews, we find issues which essentially revolve around consumers 'not knowing what they were getting into' when purchasing a device, uncovering challenges in using and understanding their new device and how it fits into their smart home environment.

2.3 Interventions in consumer IoT purchase decisions

Prior interviews with retail customers have highlighted the point of purchase as a critical point for informing decisions about the security of new computing devices [37]. In making decisions about security, home users get their security and privacy advice from various sources, where this can include family, friends, and peers [38, 39], informal technical experts [40, 41], and media such as news stories, blogs, and TV [42, 43].

Studies have analysed the marketplace for IoT devices to explore methods and mediums to inform consumers of the privacy posture of devices and their corresponding consequences. For instance, Gopavaram et al. [44] investigated customers' Willingness-To-Pay (WTP) for privacy vs their Willingness-To-Accept (WTA) a lack thereof through an emulated marketplace study. WTA participants, presented with the highest privacy settings by default, were more likely to pay a premium and purchase devices with a higher privacy rating, thereby indicating that interface design also influences purchase decisions.

Along similar lines, Blythe, Johnson and Manning [8] note that providing people with simple security-related information prior to making their purchase decision has the potential to encourage the purchase of devices which are more secure. They also found that consumers are willing to pay more for increased security and that the relative amount of risk reduction has no significant impact on that willingness.

In addition, various forms of S&P information labels have been proposed, including a graded label, labels indicating S&P features, and labels indicating 'approval' by independent assessment [45]. Blythe et al. [45] found that except for cases where an information label indicated that a device

has poor security, consumers were significantly more likely to buy a device with a label. They also found that although functionality was generally more valued, people were willing to pay the same premium for both improved functionality and security. Emami-Naeini et al. [16] presented participants with labels which were a mix of the aforementioned label types and also reported positive feedback from participants who indicated that they struggled to find this kind of information at the time of device purchase.

With a focus on the availability and regularity of security updates for IoT devices, Morgner et al. [14] found support among their survey participants for this kind of information, more so for those who perceived higher risks in using such devices. Broader efforts to both improve and standardise IoT S&P features include nation-level codes of practice (including in the UK [46, 47] and US [48, 49]), where various challenges to such device standardisation efforts have been highlighted in research [50, 51], including agreement on standards and evidencing their effectiveness. Many of these interventions seek to standardise various assurances, either from manufacturers themselves or from independent experts, that the S&P properties of a device are sufficient to be able to use a newly-purchased home IoT device securely.

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

3 Methodology

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

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

3.1 Selection of IoT devices

Research on IoT malware, botnets and compromised devices has identified specific products being compromised at scale because of severe security failures, such as using known factory-default credentials. Many of these devices fall into four categories [21]: surveillance systems (including DVR/NVR), set-top boxes, smart home hubs, and routers.

Although routers are often not considered as being an IoT device, they are integral to home networks and also susceptible to IoT-related attacks.

We approached device selection in two ways. First, we searched for specific products known to be vulnerable to IoT malware infections (specifically Mirai) [21]. This produces a set of devices with a high likelihood of prior, if not still current, security issues. We were interested to see if the reviews for these products would contain more comments on security or privacy (S&P) than other products for the same device type. We searched for these once-vulnerable devices on the Amazon websites using a combination of manufacturer, model name/number and device type (e.g., XXX YYY-123 Router, device list from Appendix A of [21]) and collected the product links. Of the 53 IoT devices we searched for, only 16 were still being sold: 14 routers, one DVR, one set-top box, and no smart home hub. The DVR and set-top box did not contain any reviews and were dropped, resulting in a single category of 14 once-vulnerable routers.

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

3.2 Product page and review retrieval

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

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

Before scraping the reviews for these products, we first reviewed all product links manually. More than 60% of them were dropped because they did not point to an IoT product. For instance, across IP cameras, there were multiple results for fake cameras that merely act as a deterrent for burglars and cameras that do not connect to the internet. With a focus on internet-enabled devices, we excluded products that do not connect to the internet or only use their proprietary

mesh network for connectivity. Devices with optional internet connectivity, like IP cameras that could be connected to the internet using a sim card, were kept in the set. Likewise, results for devices like baby monitors, spy cams, and pet cams that were internet enabled were retained. The final set consisted of 1415 product links. The count of product links collected per Amazon website and device type is Table 1.

Table 1: Count of collected product links for each device type after cleanup.

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

3.3 Review dataset construction

Using the product links we collected, we scraped the customer reviews for each product using Python scripts. For each product link, we collected two sets of reviews—one set for each research question (Figure 1). Our first research question is ‘What fraction of customer reviews for IoT devices articulate security or privacy issues?’. To answer this, we scraped the first 30 reviews (three pages) from each of the five star ratings for each product link (from one star to five stars). This resulted in 150 reviews for most products, though for some products, there were less than 30 reviews with a particular star rating, leading to a slightly smaller set.

The upper limit of 30 reviews or three pages is in line with the results of a market research study [52] that showed that only 8% of consumers read more than 26 reviews before purchase. We chose this to reflect the number of reviews a realistic, motivated buyer would read before making a product purchase. Since other studies [53] show that, on average, seven reviews are read before purchase, we found this approach better suited to study what S&P-issues a prospective buyer might encounter than collecting all product reviews.

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

Table 2: Count of reviews collected for each research question.

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

Our second research question is ‘When security or privacy issues are mentioned, what themes are being articulated?’. For this question, we wanted to focus on reviews that explicitly mention security and privacy issues. Since the top reviews collected to answer the first research question are not representative of S&P-issues raised in reviews, we used Amazon’s ‘search customer reviews’ option for each product link. The keywords used for searching were informed by prior user studies [32–36, 54]. This was done to account for how users articulate these concerns rather than the ‘tech-savvy’ words used within the research community (e.g., a user may say ‘setup’ instead of ‘configuration’). Even though terms like “get into”, “always listening” and “big data” were used by users in those studies, we did not include them in our queries as Amazon does not support concatenated search terms. The list of keywords is presented in Appendix A.2, grouped by category. For all scraped reviews, we collected the title, content, date the review was posted, the country it was posted from and the number of people that voted the review helpful. The username was not collected because of privacy considerations. The number of reviews collected for each research question across each website and device type is shown in Table 2.

3.4 Quantifying presence of IoT S&P issues

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

We used Latent Dirichlet Allocation (LDA) topic modelling to discover ‘topics’ in the review dataset [55]. Two user-defined parameters are entered into the LDA model - the number of topics (k) and the number of words for each topic. Once a set of topics has been generated, a coherence test can be used to assess the quality of the results based on the

distance between words in the same topic. However, since it is difficult to assess ‘ k ’ a priori, LDA was run for different ‘ k ’ values, and the value with the best coherence score was chosen for each model. In total, the models identified 22 topics—12 for Surveillance Systems, six for routers and two each for the remaining device types.

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

Two researchers manually labelled these 1700 reviews according to whether they discussed security or privacy-related issues or not. This activity followed a thematic analysis approach [56]. More precisely, we used ‘codebook’ thematic analysis to manually categorise the content as S&P-related or not; and inter-rater reliability does not impact the quality of these results [57, 58]. However, disagreements (7.2% for reviews from each topic and 6% for geographical comparison) in the classification were resolved through discussion and clarification of what was within the scope of S&P. In addition to obvious statements about the security or privacy of a product, reviews which referred to, e.g., firmware updates and those with mentions of authentication during setup, were also considered within scope.

3.5 Determining context of IoT S&P themes

After quantifying the presence of security and privacy (S&P) issues, we examined the themes articulated in these reviews. For this purpose, we collected 15,174 reviews via search queries with security and privacy-related keywords. We then ran the topic modelling technique to identify clusters from which we could sample reviews for qualitative thematic analy-

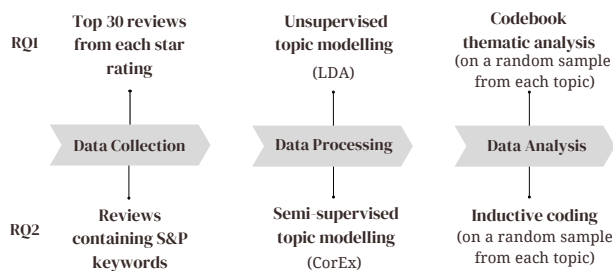


Figure 1: Overview of steps followed for each research question.

sis. For this, we chose semi-supervised Anchored Correlation Explanation (CorEx) topic modelling [59]. Unlike LDA, the CorEx topic model makes few assumptions about the latent structure of the data and flexibly incorporates domain knowledge through the anchor words that are fed to the model.

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

In the next step, a random sample of 100 reviews from each of these topic clusters was taken for thematic analysis. This sampling technique ensured that the samples taken for the thematic analysis were representative of each cluster. The thematic analysis allowed us to better understand the contexts in which S&P feature in customer reviews. The methodology outlined by Braun and Clarke [56] was followed for the thematic analysis for the context of S&P issues. One coder analysed the produced dataset, performing inductive coding to identify themes emerging from the review content [57]. Regular review meetings were held with other researchers in the team to clarify and discuss the codes, which helped trim, extend or change codes as needed. Figure 1 provides an overview of the steps followed for each research question, from data collection to analysis.

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

3.6 Research ethics

The Ethics Review Board of our institution approved the study design and data management protocol. We evaluated our research design against the principles of the Menlo Report for ethical practices in computing studies [60]. Data was collected on publicly available customer ratings and reviews on Amazon websites, and the associated usernames were not collected. Moreover, the scraping process was distributed over a longer duration through added delays in the script to ease the load on Amazon servers. Further, the scripts were fed specific pages and were not crawlers. With respect to 'justice', our study design aims to contribute to reducing asymmetry for all consumers, not specific groups.

4 Results

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

4.1 S&P prevalence in reviews

As outlined in Section 3.4, we manually analysed 50 random reviews from each of the 22 topics output by the LDA models to answer our first research question. An overview of these topics is added to Appendix A.3. We thus analysed 600 reviews for Surveillance Systems, 300 for routers, and 100 each for Hubs, Set-top boxes and Once-vulnerable Routers.

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

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

In the end, classification resulted in varying numbers of reviews per topic sample referring to S&P. The highest references were within the topics for routers (15/50). The lowest (5.9% for Surveillance Systems) is still higher than other results [30]; we revisit this discrepancy in the Discussion. On average, each sample had five reviews relating to S&P. This shows that S&P issues are not siloed in specific conversations and instead emerge in most contexts. We explore these contexts further in the next subsection.

The overall results of the classification for each device type category are shown in Table 3. Across all top reviews, about one in ten reviews (9.8%) articulate S&P-related issues. On the one hand, this is a minor fraction of all reviews. On the other hand, it does mean that potential buyers browsing reviews stand a decent chance of encountering comments on the S&P properties of the devices they are looking at. It also means that review writers feel these aspects are important enough to mention them one in ten times, which is a non-trivial amount given the brevity of most reviews: the average review in this dataset contains 100 words (median: 65). That is less than the length of this paragraph (123 words).

Table 3: Percentage of reviews referring to security and privacy issues for each device type.

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

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

4.1.1 S&P concerns in reviews across geographies

In order to compare S&P concerns across the six countries in our dataset, we drew a random sample of 100 reviews from each country and manually classified them as S&P related or not. In some cases, when there are not enough reviews for a product on the country’s website, Amazon posts reviews from other countries. However, for this analysis, the reviews were chosen based on the country it was posted from rather than the Amazon website from which it was scraped. The results are presented in Table 4. The average across all the countries (9.5%) is in line with the overall results from the device-wise analysis (9.8%), and the standard deviation is 3.67%. The deviation can be explained by the results for the US, which is higher than average (16%) and India, which is half the average (5%). For Australia and Canada, the results are in line with the global average of 10%, while for Singapore and UK, the percentage is 2% less.

Table 4: Comparison of S&P concerns across the stores.

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

4.1.2 Correlation between rating assigned and mention of S&P

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

4.2 Inductive thematic analysis results

To answer the second research question, ‘What themes are being articulated in reviews with S&P comments?’, 100 reviews were sampled randomly from each of the eight topics output by the semi-supervised algorithm (see Section 3.5). Thematic analysis was then conducted on these 800 reviews. The first step in the thematic analysis involved inductively defining and adding a single code to reviews based on their content.

We defined 99 granular codes before reaching saturation. Of the 800 reviews we analysed, 485 (60.6%) did not contain any references to security or privacy, even though they contained one of our search terms. This was an expected side-effect of using organic terms of real users, such as “record”, “log”, etc. This inevitably selects many reviews that are not related to security or privacy. These reviews did not receive a code and were not considered for further analysis. In the final step, we condensed the 99 codes into a smaller set of themes based on high-level commonalities among the codes. Appendix A.5 contains the full list of codes and themes.

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

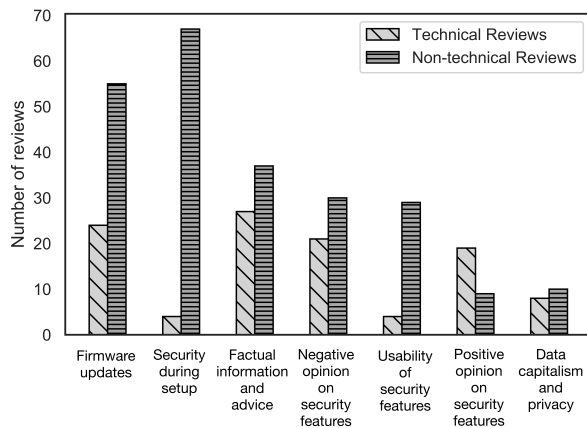


Figure 2: Distribution of technical and non-technical reviews over the seven substantive themes.

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

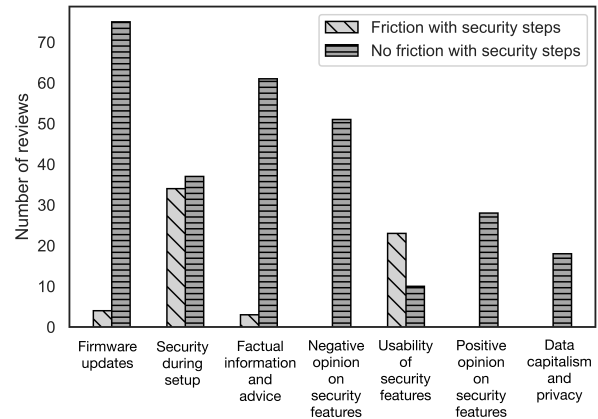


Figure 3: Distribution of reviews mentioning friction and those not over the seven substantive themes.

We added two additional codes to all reviews to complement the thematic code: whether they contained technical statements (yes or no) and whether they expressed friction with the security features (yes or no). The distribution of these two distinct themes over the seven themes is shown in Figures 2 and 3. Of the 63 reviews (20%) that express friction with security steps, none were written in technical language. Although the technical reviews also mention experiencing trouble with security features, they tend to articulate problems as specific technical critiques of security features – one of our thematic labels. In contrast, reviews in non-technical language express the trouble as personal frustration for not being able to achieve the desired result.

Table 5: Distribution of reviews over themes, friction, technical nature, and security-related versus not security-related, as well as average review score per subset.

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

Table 5 presents the distribution of reviews across the seven substantive themes and two distinctions of friction and technical nature. It also includes the average star rating of the reviews in each subset.

4.2.1 Firmware updates

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

As discussed in Section 4.1, we included all reviews discussing firmware, given its crucial role in device security and its potential relevance to prospective buyers. Of the 78 reviews in this theme, 23 talk about firmware updates in relation to issues with the device—complaints about updates not solving an issue, causing it or not helping with it. A couple of reviews mention being annoyed with frequent updates “*too many interruptions for firmware patches*” (R11035-R), while on the other hand, a handful are appreciative of it, such as “*Pros: incredibly fast, frequent updates*” (R8927-B). There were also a few reviews about firmware being “*too buggy even after updating*” (R13798-R), firmware update process being “*unnecessarily kludgy*” (R11624-R) and firmware updates containing feature updates as well. Most of the reviews about issues during firmware update mention reaching out to customer service for assistance, but without always being able to resolve them:

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

In assessing the usability of the update process itself for consumer IoT devices, Haney & Furman [61] found that participants experienced a lack of transparency in how updates worked and how. Where they noted a disconnect between updates and security, here we see a similar disconnect with firmware updates, specifically, around expectations of their role in resolving issues with devices.

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

“Was looking for a cheap router with updated Firmware available that included KRACK Patch

... I had to manually download and update the firmware...” (R11045-R)

A similar desire for timely updates has been reported elsewhere in an exploration of user information expectations for product ‘security labels’ [14].

4.2.2 Security during setup

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

Of the 172 reviews, 71 of these explicitly refer to security, e.g., “*enter a name for the SSID and create a password, and that was it*” (R11280-R). The rest do not refer to security steps like setting up usernames and passwords as an explicit step in setup, e.g., “*Easy to setup and configure*” (R11107-R). Interestingly, almost all reviews that refer to security steps during setup are non-technical reviews (see Figure 2).

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

4.2.3 Factual information and advice

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

Nine reviews outline potential security issues in devices and contain technical advice on overcoming them. This is akin to ‘informal technical support’ normally provided by a ‘local expert’ to a device user [40, 41]. Four of these reviews ask users to isolate the devices on their network:

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

Of the reviews that contain security advice in non-technical language, two advise users to change the password of their devices *“I’m sure it’s a default manufacturer password so I changed it, as I suggest anyone do” (R9144-R)*. Three warn users against buying products due to associated S&P issues:

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

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

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

4.2.4 Negative opinion on security features

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

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

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

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

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

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

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

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

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

Of note here is that some reviews situate S&P issues, where such ‘stories’ about security or privacy experiences have been seen to resonate with technology users in similar contexts [39], where stories are typically of perceived security incidents.

4.2.5 Usability of security features

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

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

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

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

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

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

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

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

4.2.6 Positive opinion about security features

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

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

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

“...Personally I’m running super secure WiFi behind an awesome firewall and a VPN, there is

noway some “hacker” is going to try that hard to get into this camera...” (R4409-S)

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

4.2.7 Data capitalism and privacy

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

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

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

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

Research looking specifically at smart speakers [32] has found that users have insufficient understanding of the data that is collected and processed by these IoT devices. In this context, this translates to reviews demonstrating concerns about having insufficient information about privacy-related matters to have made an informed purchase decision. Prior research examining reviews [30] has surfaced consumer concerns about service practices, where we evidence specific features and stages in device use where consumers focus those concerns.

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

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

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

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

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

5 Discussion

The overarching result for RQ1 across all reviews analysed shows that, on average, 9.8% of reviews refer to S&P. There is some variance across device types, from surveillance systems with 5.9% to routers with 16.3%. Interestingly, once-vulnerable routers have a slightly lower percentage (13.6%), so the security problems that have plagued these devices have not emerged in Amazon reviews. Overall, it illustrates that a notable portion of reviews discuss issues related to S&P for these device types, compared to the 2% for home assistants [30]. This difference could be from our methodology - manual analysis allows for a more nuanced and reliable classification - or our device selection strategy; the market is relatively more mature for our devices than smart home assistants. Similar to S&P considerations historically being pushed to the end of the development life cycle or treated as an afterthought [62], S&P issues might permeate the market relatively late in the product diffusion curve [63].

Further, the results from the thematic analysis for RQ2 show that customers talk about their negative personal experiences with these devices, as evidenced by the reviews that discuss frustrations with device setup, stories of hacking, and exhaustion with what is perceived as unwarranted data collection from companies. Moreover, we also see reviews providing security advice and voicing privacy concerns. Thus, our results indicate that at least a small percentage of consumers in the market for IoT devices *do* care about S&P and are able to articulate where these concerns arise in the life-cycle [64] of device ownership. There are then 'signals' in the market, the information within these reviews, that can be leveraged to ensure that other prospective buyers can make more informed decisions and to indicate to manufacturers that

consumers care about S&P. This also signifies that amplifying this information would be useful on both these counts.

Looking further at RQ2, we find that some reviewers express concerns which inform a negative view of the device overall - their personal preferences were not met. This represents *dissatisfaction* with a newly-purchased device. Others express a range of positive and negative sentiments toward the setup phase for a newly-purchased device, where the most friction with security steps was experienced during setup (Fig 3). This reveals the setup phase as a critical point in the ownership of a smart home device where information is needed.

The experiences of reviewers may point to a discounting of S&P risks in the market itself. It may be that a customer cannot return/exchange a device solely on the grounds of their S&P preferences not being met; we saw many reviews of owners being disappointed after purchase when new information about a device's S&P properties becomes apparent to them, after some not-inconsequential period of attempted use; arguably those preferences are not being treated seriously at present, outside of regular no-questions-asked return policies. However, this points to a challenge for prospective buyers to *find* that information within reviews which may also discuss non-S&P issues and preferences. Even when they are able to find these reviews, the technical language used across most of the themes (Figure 2) could hinder the usefulness of the information for those without a technical background.

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

5.1 Limitations

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

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

Our analysis is limited to product reviews on an online marketplace. It is acknowledged that the potential for vari-

ous forms of media to inform security behaviours is under-researched (including TV and news [43, 65]); we add product reviews to that list, where here we found potential to leverage this additional source of information. Many of the reviews we analysed focused on the initial experience of device ownership, which presumably influenced the time when most users write reviews, that is, right after purchase. Where other research has recently begun to detail milestones in extended IoT device use (e.g., [66]), our analysis here pinpoints specific activities where S&P-issues were found, such as device setup.

5.2 Recommendations

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

- **Highlight S&P-related reviews.** Given that a not-inconsiderable portion of the reviewer population (9.8%, as in Table 3) articulates S&P-related concerns, it would be useful if e-commerce websites highlight these reviews. Irrespective of whether the reviews are fake or misinformed, highlighting the S&P signals in such reviews might encourage other buyers browsing through reviews to factor S&P prior to making a purchase decision. However, highlighting S&P reviews will only be partially successful since many of these (33.7%) are technical in nature and language. In order for it to be useful for a less knowledgeable buyer (Section 4.2.3), it would help to have an S&P specific rating for the devices. Such a rating will serve as a shorthand indicator of device quality and help consumers from a non-technical background interpret the sentiment expressed in these reviews.
- **Use the review system to match advice to emergent concerns.** Consumers may be willing to follow S&P advice if it addresses their existing concerns. The negative opinions examined in Section 4.2.4 surfaced S&P concerns and stories, for instance, and in Section 4.2.5, unusable security features were seen to hinder device use. If support can be provided to configure these features to the satisfaction of the user, it could reduce the lack of engagement with the devices; results in Section 4.2.6 highlight that some reviewers were able to find appreciable S&P features. On online shopping platforms such as Amazon, this could be addressed in the Q&A part of a product listing, with a variation of ‘signals of interest’ [12]. The market may not appreciate that a device has been bought but *later* unused due to S&P concerns – a further indication of ‘interest’ in using the device is then necessary to indicate that a solution was subsequently found. Where answers for S&P-related questions on the shopping platform can be indicated as useful, it can ‘match’ a solution to a concern and restore intentions to use a product beyond the initial signal of device purchase.

- **Design for shortcuts in security features.** We noticed a distinction between reviews framed in technical details, and others not (Section 4.2.2 and Figure 2), with similar concerns around shared device use activities (such as setup, adding a new user, etc.). There is then scope to balance the needs of users who want detailed configuration options and novice users who want the ease of setup. This relates to the established design principle of “Flexibility and efficiency of use” [67], and provides configuration options for both experienced and inexperienced users.

6 Conclusions

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

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

Acknowledgments

This publication is part of the RAPID project (Grant No. CS.007) financed by the Dutch Research Council (NWO). The authors would like to thank Elsa Turcios Rodríguez for her valuable feedback during the early work on this paper.

References

- [1] M. Antonakakis, T. April, M. Bailey, E. Bursztein, J. Cochran, Z. Durumeric, J. Alex Halderman, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, Y. Zhou, M. Antonakakis Tim April, M. Bernhard Elie Bursztein, J. J. Cochran Zakir Durumeric Alex Halderman Luca Invernizzi, M. Kallitsis, D. Kumar, C. Lever Zane Ma, J. Mason, and N. Sullivan Kurt Thomas, “Understanding the Mirai Botnet,” *USENIX Security ‘17*, 2017. [Online]. Available:

<https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis>

- [2] C. Koliass, G. Kambourakis, A. Stavrou, and J. Voas, “DDoS in the IoT: Mirai and other botnets,” *Computer*, vol. 50, no. 7, pp. 80–84, 2017.
- [3] K. Thomas, D. Huang, D. Wang, E. Bursztein, C. Grier, T. J. Holt, C. Kruegel, D. McCoy, S. Savage, and G. Vigna, “Framing dependencies introduced by underground commoditization,” *Workshop on the Economics of Information Security (WEIS)*, 2015.
- [4] H. L. Bijmans, T. M. Booij, and C. Doerr, “Just the tip of the iceberg: Internet-scale exploitation of routers for cryptojacking,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 449–464.
- [5] N. Bowles, “Thermostats, locks and lights: Digital tools of domestic abuse,” *The New York Times*, vol. 23, 2018.
- [6] L. M. Tanczer, I. López-Neira, and S. Parkin, “‘I feel like we’re really behind the game’: Perspectives of the United Kingdom’s intimate partner violence support sector on the rise of technology-facilitated abuse,” *Journal of Gender-Based Violence*, 2021.
- [7] B. Schneier, *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company, 2018.
- [8] J. M. Blythe, S. D. Johnson, and M. Manning, “What is security worth to consumers? Investigating willingness to pay for secure Internet of Things devices,” *Crime Science*, vol. 9, no. 1, pp. 1–9, 2020.
- [9] S. Gopavaram, J. Dev, S. Das, and L. J. Camp, “IoT Marketplace: Willingness-To-Pay vs. Willingness-To-Accept,” in *Proceedings of the 20th Annual Workshop on the Economics of Information Security (WEIS 2021)*, 2021.
- [10] G. A. Akerlof, “The market for “lemons”: Quality uncertainty and the market mechanism,” in *Uncertainty in economics*. Elsevier, 1978, pp. 235–251.
- [11] M. Spence, “Job market signaling,” *The Quarterly Journal of Economics*, vol. 87, no. 3, pp. 355–374, 1973. [Online]. Available: <http://www.jstor.org/stable/1882010>
- [12] A. E. Roth, *Who gets what—and why: The new economics of matchmaking and market design*. Houghton Mifflin Harcourt, 2015.
- [13] M. B. Akdeniz, R. J. Calantone, and C. M. Voorhees, “Signaling quality: An examination of the effects of marketing-and non marketing-controlled signals on perceptions of automotive brand quality,” *Journal of Product Innovation Management*, vol. 31, no. 4, pp. 728–743, 2014.
- [14] P. Morgner, C. Mai, N. Koschate-Fischer, F. Freiling, and Z. Benenson, “Security update labels: Establishing economic incentives for security patching of IoT consumer products,” in *2020 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2020, pp. 429–446.
- [15] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, “Ask the experts: What should be on an IoT privacy and security label?” in *2020 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2020, pp. 447–464.
- [16] P. Emami-Naeini, H. Dixon, Y. Agarwal, and L. F. 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*, 2019, pp. 1–12.
- [17] V. Garg, “A lemon by any other label.” *ICISSP*, pp. 558–565, 2021.
- [18] F. Zhu and X. Zhang, “Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics,” *Journal of marketing*, vol. 74, no. 2, pp. 133–148, 2010.
- [19] D. K. Basdeo, K. G. Smith, C. M. Grimm, V. P. Rindova, and P. J. Derfus, “The impact of market actions on firm reputation,” *Strategic Management Journal*, vol. 27, no. 12, pp. 1205–1219, 2006. [Online]. Available: <http://www.jstor.org/stable/20142408>
- [20] S. He, B. Hollenbeck, and D. Proserpio, “The market for fake reviews,” *Marketing Science*, 2022.
- [21] E. Rodríguez, A. Noroozian, M. van Eeten, and C. Gañán, “Superspreaders: Quantifying the role of IoT manufacturers in device infections,” in *20th Workshop on the Economics of Information Security (WEIS)*, 2021.
- [22] S.-H. Lee, “How do online reviews affect purchasing intention?” *African Journal of Business Management*, vol. 3, no. 10, pp. 576–581, 2009.
- [23] J. Zhan, H. T. Loh, and Y. Liu, “Gather customer concerns from online product reviews—a text summarization approach,” *Expert Systems with Applications*, vol. 36, no. 2, pp. 2107–2115, 2009.
- [24] D.-H. Park, J. Lee, and I. Han, “The effect of online consumer reviews on consumer purchasing intention: The moderating role of involvement,” *International journal of electronic commerce*, vol. 11, no. 4, pp. 125–148, 2007.

- [25] C.-L. Lin, S.-H. Lee, and D.-J. Horng, “The effects of online reviews on purchasing intention: The moderating role of need for cognition,” *Social Behavior and Personality: an international journal*, vol. 39, no. 1, pp. 71–81, 2011.
- [26] N. Hu, P. A. Pavlou, and J. J. Zhang, “On self-selection biases in online product reviews,” *MIS Q.*, vol. 41, no. 2, pp. 449–471, 2017.
- [27] S. Han and C. K. Anderson, “Customer motivation and response bias in online reviews,” *Cornell Hospitality Quarterly*, vol. 61, no. 2, pp. 142–153, 2020.
- [28] M. Trenz and B. Berger, “Analyzing online customer reviews - An interdisciplinary literature review and research agenda,” 2013.
- [29] L. Manikonda, A. Deotale, and S. Kambhampati, “What’s up with privacy? User preferences and privacy concerns in intelligent personal assistants,” in *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 2018, pp. 229–235.
- [30] N. Fruchter and I. Liccardi, “Consumer attitudes towards privacy and security in home assistants,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1–6. [Online]. Available: <https://doi.org/10.1145/3170427.3188448>
- [31] D. van der Linden, M. Edwards, I. Hadar, and A. Zamansky, “Pets without PETS: On pet owners’ underestimation of privacy concerns in pet wearables,” *Proc. Priv. Enhancing Technol.*, vol. 2020, no. 1, pp. 143–164, 2020.
- [32] N. Abdi, K. M. Ramokapane, and J. 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)*, 2019, pp. 451–466.
- [33] E. Zeng, S. Mare, and F. Roesner, “End user security and privacy concerns with smart homes,” in *thirteenth symposium on usable privacy and security (SOUPS 2017)*, 2017, pp. 65–80.
- [34] M. Williams, J. R. Nurse, and S. Creese, “Privacy is the boring bit: user perceptions and behaviour in the Internet-of-Things,” in *2017 15th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 2017, pp. 181–18 109.
- [35] S. Zheng, N. Apthorpe, M. Chetty, and N. Feamster, “User perceptions of smart home IoT privacy,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, pp. 1–20, 2018.
- [36] J. M. Haney, S. M. Furman, and Y. Acar, “Smart home security and privacy mitigations: Consumer perceptions, practices, and challenges,” in *International Conference on Human-Computer Interaction*. Springer, 2020, pp. 393–411.
- [37] S. Parkin, E. M. Redmiles, L. Coventry, and M. A. 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.
- [38] S. Das, L. A. Dabbish, and J. I. Hong, “A typology of perceived triggers for end-user security and privacy behaviors,” in *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*, 2019, pp. 97–115.
- [39] E. Rader, R. Wash, and B. Brooks, “Stories as informal lessons about security,” in *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS)*, 2012, pp. 1–17.
- [40] E. S. Poole, M. Chetty, T. Morgan, R. E. Grinter, and W. K. Edwards, “Computer help at home: methods and motivations for informal technical support,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2009, pp. 739–748.
- [41] N. Nthala and I. Flechais, “Informal support networks: An investigation into home data security practices,” in *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*, 2018, pp. 63–82.
- [42] S. Das, J. Lo, L. Dabbish, and J. 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*, 2018, pp. 1–12.
- [43] E. Rader and R. Wash, “Identifying patterns in informal sources of security information,” *Journal of Cybersecurity*, vol. 1, no. 1, pp. 121–144, 2015.
- [44] S. R. Gopavaram, J. Dev, S. Das, and J. Camp, “IoT-Marketplace: Informing purchase decisions with risk communication,” 2019.
- [45] S. D. Johnson, J. M. Blythe, M. Manning, and G. T. Wong, “The impact of IoT security labelling on consumer product choice and willingness to pay,” *PloS one*, vol. 15, no. 1, p. e0227800, 2020.
- [46] UK Department for Digital, Culture, Media & Sport (DCMS), “Code of practice for consumer IoT security,” <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security>, UK Department for Digital, Culture, Media & Sport (DCMS), 2018.

- [47] —, “Regulating consumer smart product cyber security - government response,” <https://www.gov.uk/government/publications/regulating-consumer-smart-product-cyber-security-government-response>, UK Department for Digital, Culture, Media & Sport (DCMS), 2021.
- [48] M. Fagan, M. Yang, A. Tan, L. Randolph, and K. Scarfone, “Security review of consumer home Internet of Things (IoT) products,” <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8267-draft.pdf>, US National Institute of Standards and Technology (NIST), 2019.
- [49] K. Megas, B. Cuthill, and S. Gupta, “Establishing confidence in iot device security: How do we get there?(draft),” National Institute of Standards and Technology, Tech. Rep., 2021.
- [50] S. L. Keoh, S. S. Kumar, and H. Tschofenig, “Securing the Internet of things: A standardization perspective,” *IEEE Internet of things Journal*, vol. 1, no. 3, pp. 265–275, 2014.
- [51] E. Leverett, R. Clayton, and R. Anderson, “Standardisation and certification of the ‘Internet of Things’,” in *Proceedings of the Workshop on Economics of Information Security (WEIS)*, vol. 2017, 2017.
- [52] D. McMillen, “Internet of Threats: IoT Botnets drive surge in network attacks,” Apr 2021. [Online]. Available: <https://securityintelligence.com/posts/internet-of-threats-iot-botnets-network-attacks/>
- [53] J. Fogel and S. Zachariah, “Intentions to use the yelp review website and purchase behavior after reading reviews,” *Journal of theoretical and applied electronic commerce research*, vol. 12, no. 1, pp. 53–67, 2017.
- [54] C. Geeng and F. Roesner, “Who’s in control? Interactions in multi-user smart homes,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–13.
- [55] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent Dirichlet Allocation,” *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [56] V. Clarke and V. Braun, *Thematic Analysis*. New York, NY: Springer New York, 2014, pp. 1947–1952. [Online]. Available: https://doi.org/10.1007/978-1-4614-5583-7_311
- [57] V. Braun and V. Clarke, “One size fits all? What counts as quality practice in (reflexive) thematic analysis?” *Qualitative research in psychology*, pp. 1–25, 2020.
- [58] N. McDonald, S. Schoenebeck, and A. Forte, “Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, pp. 1–23, 2019.
- [59] R. J. Gallagher, K. Reing, D. Kale, and G. Ver Steeg, “Anchored correlation explanation: Topic modeling with minimal domain knowledge,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 529–542, 2017.
- [60] E. Kenneally and D. Dittrich, “The Menlo report: Ethical principles guiding information and communication technology research,” *Available at SSRN 2445102*, 2012.
- [61] J. M. Haney and S. M. Furman, “Work in progress: Towards usable updates for smart home devices,” in *International Workshop on Socio-Technical Aspects in Security and Trust*. Springer, 2020, pp. 107–117.
- [62] C. Steward Jr, L. A. Wahsheh, A. Ahmad, J. M. Graham, C. V. Hinds, A. T. Williams, and S. J. DeLoatch, “Software security: The dangerous afterthought,” in *2012 Ninth International Conference on Information Technology-New Generations*. IEEE, 2012, pp. 815–818.
- [63] E. M. Rogers, A. Singhal, and M. M. Quinlan, “Diffusion of innovations,” in *An integrated approach to communication theory and research*. Routledge, 2014, pp. 432–448.
- [64] D. Kotz and T. Peters, “Challenges to ensuring human safety throughout the life-cycle of smart environments,” in *Proceedings of the 1st ACM Workshop on the Internet of Safe Things*, 2017, pp. 1–7.
- [65] E. M. Redmiles, S. Kross, and M. L. Mazurek, “How I learned to be secure: A census-representative survey of security advice sources and behavior,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 666–677.
- [66] G. Chalhoub, M. J. Kraemer, N. Nthala, and I. 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*, 2021, pp. 1–16.
- [67] J. Nielsen, *Usability engineering*. Morgan Kaufmann, 1994.

A Appendix

A.1 Search terms used for each of the four device types

Surveillance Systems: Surveillance camera, Network Camera, IP Camera, Security Camera, Dome Camera and DVR/NVR

Smart home hub: Smart home hub, Smart home control panel, Smart home automation system

Set-top box: Digital set top box, IP set top box

Router: Router, WiFi repeater

A.2 Search terms used for identifying S&P related customer reviews by category

Configuration and Authentication: setup, configure, profile, default, control, 2FA, authentication, password

Access and Storage: access, manipulate, watch, track, record, log, data

Encryption and Security: encrypt, encryption, ssl, protocol, secure

Privacy: privacy, private, personal, trust, open

Attack: hack, attack, fraudsters, spy, steal, blackmail, criminals, cutoff

Patches and updates: patch, uptodate, update, firmware, vulnerability, risk, safe, protect

A.3 Results of LDA

LDA for Surveillance Systems

Topic 1: night, vision, product, quality, picture, light, day, time, image, work || **Topic 2:** app, phone, time, monitor, quality, video, home, view, picture, work || **Topic 3:** system, quality, nvr, image, security, poe, setup, setting, video, picture || **Topic 4:** network, device, connection, work, setup, router, app, internet, issue, access || **Topic 5:** unit, battery, ring, model, number, year, doorbell, resolution, zone, month || **Topic 6:** cloud, storage, subscription, video, window, stream, service, option, plan, year || **Topic 7:** motion, detection, notification, alert, time, sensitivity, record, video, setting, alarm || **Topic 8:** cam, brand, contact, color, noise, today, stuff, good, audio, condition || **Topic 9:** cable, power, wire, ethernet, wall, box, plug, screw, plastic, plate || **Topic 10:** card, video, sd, record, app, footage, recording, memory, playback, file || **Topic 11:** software, car, pc, door, computer, web, hardware, people, foot, interface || **Topic 12:** support, customer, service, issue, problem, email, help, tech, replacement, update

LDA for Routers

Topic 1: extender, unit, range, room, instruction, wifi, work, light, install, plug || **Topic 2:** network, setup, system, app, home, mesh, point, access, port, cable || **Topic 3:** connection, performance, laptop, bit, set, video, quality, eero, phone, work || **Topic 4:** speed, signal, house, internet, coverage, floor, strength, wifi, drop, test || **Topic 5:** support, time, work, day,

money, service, hour, tech, week, customer || **Topic 6:** issue, price, month, review, year, band, problem, time, model, day

LDA for Hubs

Topic 1: device, time, app, work, product, tv, control, button, setup, hub || **Topic 2:** music, speaker, sound, play, love, sound_quality, room, alarm, quality, question

LDA for Set-top Boxes

Topic 1: tv, box, fire, device, app, stick, work, product, control, issue || **Topic 2:** channel, tv, time, record, program, device, unit, guide, cable, recording

LDA for Once-vulnerable Routers

Topic 1: network, setup, connection, speed, work, range, feature, access, signal, option || **Topic 2:** issue, time, internet, connection, day, problem, cable, work, unit, support

A.4 Results of Anchored CorEx

Topic 1 (23.8%) : *anchor words:* setup, configure, control - setup, password, easy setup, setup easy, easy, initial setup, configure, username password, username, camera setup || **Topic 2 (19.8%) :** (*anchor words:* access, watch, record) - record, camera, motion, night, detection, quality, vision, night vision, motion detection, video || **Topic 3 (11.1%) :** (*anchor words:* encrypt, secure, protocol) - secure, protocol, encrypt, address, iris, blue iris, onvif || **Topic 4 (15.6%) :** (*anchor words:* open, trust, personal) - open, trust, personal, open source, time open, open camera, personal data, camera open, open door, seal || **Topic 5 (11.5%) :** (*anchor words:* steal, hack, spy) - hack, steal, price, easily, attach, small, hole, recommend, outside, white || **Topic 6 (8.9%) :** (*anchor words:* protect, update, firmware) - update, firmware, firmware update, update firmware, latest, latest firmware, version, update review, upgrade, firmware upgrade || **Topic 7 (6.5%) :** *no anchors* - support, work, time, issue, review, try, problem, email, reset, contact || **Topic 8 (2.8%) :** *no anchors* - network, connect, setting, devices, cable, power, connection, feature, point, plug

A.5 Codebook from Inductive Coding

Firmware updates: fw-security-patch, fw-auto-update, exclusive-fw-update, fw-available, fw-buggy, fw-comparison, fw-language, fw-latest, fw-update-caused-issue, fw-update-didn't-fix-issue, fw-update-difficult, fw-update-for-feature, fw-update-frequent-good, fw-update-had-feature, fw-update-kludgy, fw-update-might-cause-issue, fw-update-might-fix, fw-update-not-available, fw-update-timing, fw-update-to-solve-issue, fw-update-took-longer, fw-update-when-setting-up, fw-updates-back-to-back-annoying

Security during setup: setup-password, setup-qr-code-easy, setup-qr-code-strange, setup-ssid-not-hidden, setup-too-simple, setup-wps, setup-wps-easy

Factual information and advice: chinese-servers-distrust, contacts-mfg-server, fw-update-not-available-chinese,

modem-too-secure, remote-access-cust-support, rtsp-pw-plaintext, security-advice, security-comment, security-protocol, unconcerned-about-vulnerabilities, used-device-pwd-already-set, vulnerable-to-chinese-hackers, woods-tablet

Negative opinion on security features: 2fa-missing, guest-network-doesn't-isolate, has-security-vulnerability, I-was-hacked-stories, limited-security, remote-access-uncomfortable, security-concern, unhappy-about-security, unsecure-device-alert

Usability of security features: 2FA-useless, can't-change-password, can't-setup-security, conflict-with-wpa3-wps-settings, encrypted-affects-playback, encryption-program-marketing, extra-security-on-trial, locked-out-cos-password, no-password-good, password-not-available, password-

reset-thru-customer-support, passwords-cumbersome, security-options-at-launch, share-admin-privileges, share-password-with-qr, unhappy-about-update-for-privacy

Positive opinion of security features: 2fa-better-security, configured-for-security-so-not-worried, ddos-protection-good, firewall-configuration-not-needed, firewall-good, good-securitywise, guest-network-for-security, malicious-activity-monitoring-valuable, notification-when-new-device-joins, secure-video-sharing

Data capitalism and privacy: app-permissions, cloud-not-needed, face-recognition-good-for-privacy, not-always-listening-is-plus, privacy-concern, privacy-mode-in-camera