

Which Privacy and Security Attributes Most Impact Consumers' Risk Perception and Willingness to Purchase IoT Devices?

Pardis Emami-Naeini
University of Washington
pardis@cs.washington.edu

Janarth Dheenadhyalan
Carnegie Mellon University
janarth@cmu.edu

Yuvraj Agarwal
Carnegie Mellon University
yuvraj@cs.cmu.edu

Lorrie Faith Cranor
Carnegie Mellon University
lorrie@cmu.edu

Abstract—In prior work, researchers proposed an Internet of Things (IoT) security and privacy label akin to a food nutrition label, based on input from experts. We conducted a survey with 1,371 Mechanical Turk (MTurk) participants to test the effectiveness of each of the privacy and security attribute-value pairs proposed in that prior work along two key dimensions: ability to convey risk to consumers and impact on their willingness to purchase an IoT device. We found that the values intended to communicate increased risk were generally perceived that way by participants. For example, we found that consumers perceived more risk when a label conveyed that data would be sold to third parties than when it would not be sold at all, and that consumers were more willing to purchase devices when they knew that their data would not be retained or shared with others. However, participants' risk perception did not always align with their willingness to purchase, sometimes due to usability concerns. Based on our findings, we propose actionable recommendations on how to more effectively present privacy and security attributes on an IoT label to better communicate risk to consumers.

Index Terms—Internet of Things (IoT), Privacy and Security, Label, Risk Perception, Willingness to Purchase.

I. INTRODUCTION

Consumers are concerned about IoT device data collection, how data might be used and shared, and the lack of controls for managing device privacy and security [1]–[4]. Research has shown that consumers would like to use privacy and security information when determining which device to purchase [5], but this information is not readily available [6].

A usable IoT privacy and security label could help consumers consider privacy and security when purchasing a device and make more informed purchase decisions. Although there have been several legislative proposals for IoT labels [7]–[12], few details have been specified about what these labels should include. To the best of our knowledge, a recent paper by Emami-Naeini et al. [13] is the only comprehensive proposal for an IoT privacy and security label.

After interviewing a diverse set of privacy and security experts, Emami-Naeini et al. [13] specified 47 privacy, security, and general attributes that should be included on a two-layer IoT nutrition label. In addition, they conducted a small-scale interview study with 15 IoT consumers to observe whether participants understand the information presented on the label. Although this work is a good first step toward designing effective privacy and security labels, their proposed label (see Appendix A) was primarily based on the opinion of experts, who often perceive risks differently than the public [14].

Our primary research goal is to assess which previously-identified IoT privacy and security label attributes significantly influence risk perception and willingness to purchase and in what ways. Our secondary goal is to recommend improvements to the proposed label design based on this assessment, through identification of common misconceptions that might be corrected through alternate wording or further explanation.

To achieve our research objectives and bridge the gap between experts' knowledge and consumers' understanding, we conducted a large-scale mixed-design survey study on Amazon Mechanical Turk with 1,371 participants. We considered two types of smart devices and three device recipients, which led to a total of six experimental conditions.

In our study, we tested 16 privacy and security attributes. Out of those, 15 attributes had two possible values each, corresponding to its most protective or least protective levels, while one attribute had three possible values. For each participant, we randomly selected and presented three of the 16 privacy and security attributes and for each selected attribute, we randomly selected one of its possible values. Our method effectively adds another value to each attribute (i.e., *absent*), which corresponds to the attribute not being shown to a participant.

To be specific, we presented each participant with three randomly-selected scenarios describing a hypothetical purchase setting for an IoT device with a label containing a single privacy or security attribute-value pair and a brief explanation of the attribute-value designed for consumers (see Table III in Appendix C). We measured how each attribute-value pair would change risk perception and willingness to purchase the smart device when presented in isolation.

We found that attribute-value pairs indicating that collected information could be sold or that devices lacked access control significantly elevated perceived risks and decreased the desire to purchase the device. On the other hand, attribute-value pairs indicating that no information is being shared with third parties or retained on the cloud significantly reduced participants' risk perception and increased their willingness to purchase the device. Moreover, we found that who the device is being purchased for does not significantly impact participants' risk perception, but does impact their willingness to purchase.

Our analysis shows that while an increase in participants' perceived risk is generally aligned with a reduced willingness to purchase and vice versa, they are not always aligned perfectly. In particular, we found that privacy and security attributes are more powerful in influencing consumers' risk perception than

their willingness to purchase the device.

We make the following contributions in this paper:

- Through our quantitative data collection, we identify the privacy and security attributes and corresponding values that most impact participants' risk perception and willingness to purchase IoT devices.
- Through our qualitative data collection, we gain insights into why participants were influenced or not influenced by label attributes to perceive a device as riskier or less risky, or to report being more likely or less willing to purchase that device.
- We distill a set of actionable recommendations on how to better inform consumers' purchase behavior by more effectively conveying privacy and security risks to consumers on an IoT label.

II. BACKGROUND AND RELATED WORK

We provide background on risk perception. We then discuss factors impacting consumers' willingness to purchase. Finally, we discuss research on how labels have been used to inform consumers' purchase behavior.

A. Risk Perception

Perceived risk is a subjective assessment of the likelihood of a specific event happening and concern about its consequences [15]. Research has shown that people tend to base their decisions on the perceived risk [16], [17].

In the context of privacy and security, researchers have shown that users' lack of risk awareness and knowledge about how their data might be used [18]–[21] influences their risk judgement [22], [23]. Skirpan et al. found identity theft, account breach, and job loss as the top-three rated risk scenarios related to emerging technologies [24].

Researchers have examined how people perceive risks of smart devices. Wieneke et al. conducted a qualitative study on how privacy affects decision making related to wearable devices [25]. They found that users' lack of awareness impacts their risk perception and they also observed a disparity between risk perception and behaviors. Their findings are aligned with other work in this space [19], [26].

B. Willingness to Purchase

Willingness to purchase is an indicator of actual purchase behavior [27] and has been shown to have a high correlation with it [28]–[30]. Researchers have identified a number of factors impacting consumers' purchase behavior including price, features, perceived quality, brand, social media, word of mouth, and usability [31]–[34].

Privacy is one of the concerns people have when participating in e-commerce [35]–[37]. In a study by Tsai et al., availability of accessible privacy information in search results encouraged consumers to purchase from privacy-protective websites, despite their higher prices [38]. Similarly, Kelly et al. found that consumers will engage in more privacy-protective app-selection behaviors when concise privacy information is available [39].

C. Labels

Labels are a common approach in contexts such as food [40] and energy ratings [41], [42] to effectively communicate important information to consumers. Despite their limitations [43], [44], food nutrition labels have been shown to significantly inform consumers' purchase decisions [45], [46]. In the privacy context, Apple has recently required app developers to provide information about the privacy practices of their apps. This information is presented on a privacy label located in the app store to help users make more informed app selection. [47]

Researchers have suggested that privacy and security labels for smart devices could effectively inform consumers. Emami-Naeini et al. conducted a small-scale qualitative study to explore how IoT consumers would react to privacy and security labels for smart devices. They found that consumers are generally unable to find privacy and security information for smart devices they are considering purchasing and would be interested in having privacy and security labels readily available [5].

Policymakers [7]–[10], [12], industry groups [48], [49], and certification bodies [50] have expressed interest in privacy and security labels for IoT devices. However, there has been little discussion of label format and content. Emami-Naeini et al. took a first step toward designing an informative IoT privacy and security label by interviewing and surveying experts [13]. They specified 47 important factors and proposed a layered label to present those factors.

In this study, we focused on consumers' risk perception, which often differs from that of experts [14]. We measured the significance of IoT privacy and security attributes identified by Emami-Naeini et al. [13], along with factors previously found to explain risk perception, including risk target [51]–[53], familiarity with the technology [54], [55], and attitudes and concerns [56], [57]. Specifically, we considered the device recipient (e.g., purchase for self or for someone else) to evaluate the risk target, checked whether participants owned that type of device to gauge familiarity with the technology, and varied the type of device to gauge the impact of concerns related to the type of collected data.

III. METHOD

We conducted an online study in January 2020 on Amazon Mechanical Turk (MTurk) with 1,710 participants (reduced to 1,371 participants after filtering the responses). In this section, we discuss our study design, data analysis procedures, and limitations. The study protocol was approved by our Institutional Review Board (IRB).

A. Study Design

We designed our study with two between-subject factors—the *device type* and the *recipient of the device*. We tested two types of devices and three types of device recipients for a total of six experimental conditions. Our within-subject factor was the IoT-related privacy and security information conveyed on the label. To mitigate survey fatigue [58] and keep the completion time under 15 minutes, we *randomly* assigned each participant to answer questions about only 3 of the 33 possible

pairs of attributes and their corresponding values, all associated with one randomly assigned experimental condition. The survey questions are provided in Appendix B.

1) *Pilot Survey*: Prior to launching the main study, we piloted our survey on MTurk with 50 participants. We found that more than half of the answers to open-ended questions were irrelevant words such as “nice,” “good,” or “yes,” which suggested that bots might be answering our survey [59]. We significantly improved the quality of the responses by adding a Google reCAPTCHA question [60] at the beginning of the survey to identify and eliminate bots.

2) *Participant Recruitment and Compensation*: We recruited 1,710 MTurk Master Workers from the United States who were at least 18 years old and who had a HIT, i.e., Human Intelligence Task, approval rate of at least 90%. We introduced the survey as a study about behaviors and attitudes toward smart devices. On average, it took participants 13 minutes to answer the survey questions and we paid them each \$2.50.

3) *Survey Procedure*: After presenting participants with the consent form and CAPTCHA verification, we asked about participants’ concern level and purchase history for the smart device that was assigned to their study condition. We then presented each participant with three randomly assigned hypothetical scenarios about the purchase of a smart device, using the device type and recipient in their assigned condition. Each purchase scenario included mention of a product label with a single attribute-value pair, selected at random:

Imagine you are making a decision to purchase a [device type] for [device recipient]. This device has a [device sensor] that will [device data collection]. The price of the device is within your budget and the features are all what you would expect from a [device type]. On the package of the device, there is a label that explains the privacy and security practices of the [device type].

The label on the device indicates the following:
[attribute: value] (consumer explanation)

For each of the three scenarios, we asked participants how the information on the label would change their risk perception and their willingness to purchase, the reasons behind their assessments, and a question to check whether they were paying attention to the label. We then asked a question to capture participants’ understanding of how their assigned smart device collects data. We ended the survey with demographic questions.

4) *Between-Subject Factors*: We considered device type as a between-subject factor and tested two types of devices. We selected smart speakers (with a microphone that will listen and respond to voice commands), which we hypothesized that most participants would find concerning [26], [61], and smart light bulbs (with a presence sensor that detects whether someone is present in the room to control the lighting automatically) that we expected to be perceived as not concerning [5], [26].

Our other between-subject factor was the IoT device recipient. We were interested in understanding whether participants have different risk perceptions and desires to purchase based on whom they are purchasing the device for. Hence, we tested

three conditions: Purchasing the device for oneself, gifting it to a family member, or gifting it to a friend.

5) *Concern Level and Purchase History*: To test our hypothesis on the assessed level of concern for the two tested device types, we asked participants to specify how concerned they were about the smart device collecting data and the reason for their answer. If they currently have a smart device of that type in their home, we then asked them when and how they acquired their devices. If they did not have the smart device, we asked them whether they had ever been in the market to purchase it and if so, we asked them what made them decide not to purchase it (see Appendix B-A).

6) *Privacy and Security Label Attributes*: Emami-Naeini et al. specified 47 attributes to include on the label, of which 25 are directly related to privacy or security [13]. There were two types of attributes on their label: 14 had enumerated values and the rest had URL as their values. In our study, we included 17 privacy and security attributes. We tested the 14 proposed label attributes with enumerated values. In addition, we selected 3 sub-attributes from Emami-Naeini et al.’s proposed label [13] and combined them into an additional “control” attribute. Finally, due to the importance of security patches in IoT standards and guidelines [62]–[67], we also included an attribute related to time-to-patch that was not part of the proposed label.

Since Emami-Naeini et al. [13] did not enumerate all of the attribute-values, we synthesized the possible values each attribute might take from a review of IoT privacy and security standards and guidelines. For each attribute, we identified a value that we hypothesized to be perceived as the most protective and one that we expected to be perceived as the least protective to test in our study. For one of the attributes (user controls), we considered three values; the proposed label actually presents these as binary values associated with three separate attributes (data stored on device, data stored on cloud, and local data retention time). Out of these 33 attribute-value pairs (shown in Table I), each participant answered questions related to three randomly-selected attribute-value pairs, contextualized with a hypothetical purchase scenario. We implemented the scenario selection function so that participants’ assigned attribute-value pairs would not include multiple pairs with the same attribute. We provided a consumer-friendly explanation next to each attribute-value pair (shown in Appendix C).

7) *Questions About Each Scenario*: To evaluate how well participants believed that they understood the tested attribute-value pairs and associated explanations, we asked them how confident they were that they knew what the presented information meant (see Appendix B-B).

To understand participants’ risk perception, we asked them to specify how the presented attribute-value changes the privacy and security risks they associated with the device in question (see Appendix B-B1). We then asked participants to explain the reason behind their choice. We asked similar questions to understand the impact of the privacy and security attributes on changing participants’ willingness to purchase the device (see Appendix B-B2).

Layer	Attribute	Tested value	
		Most protective	Least protective
Primary	Security update	Automatic	None
	Access control	Multi-factor authentication	None
	Purpose	Device function	Monetization
	Device storage	None	Identified
	Cloud storage	None	Identified
	Shared with	None	Third parties
	Sold to	None	Third parties
	Control over	Cloud data deletion Device storage	
Secondary	Average time to patch	1 month	6 months
	Security audit	Internal & external	None
	Collection frequency	On user demand	Continuous
	Sharing frequency	On user demand	Continuous
	Device retention	None	Indefinite
	Cloud retention	None	Indefinite
	Data linkage	None	Internal & external
	Inference	None	Characteristics and psychological traits, attitudes and preferences, aptitudes and abilities, and behaviors
	Control over	Device retention	

TABLE I: The 16 security and privacy attributes along with the values of each attribute tested. The attributes are grouped here according to the layers proposed in [13]. Note that the attribute “control over” is on both layers. Attribute “average time to patch” was not included in the proposed label. Due to the importance of security patches, we hypothesized that this attribute might be appropriate for the secondary layer.

To evaluate participants’ attention and understanding of the presented label information, we tested participants on the specific privacy and security information they were asked about with a multiple-choice question. For example, if the presented attribute-value was *security audit: none*, we asked “Which statement is correct about the device described in the previous question?” and provided three incorrect answers alongside the correct answer “The manufacturer does not conduct security audits on its devices and services.” We designed these questions to be answerable with only the information we provided in the consumer explanation, e.g., in this example, without knowing anything about the implications of security audits.

8) *Perceived Device Functionality*: To understand how participants perceived the device data collection, we asked them whether they believed the device is always sensing, sensing only when it is triggered (e.g., by mentioning the wake word or by someone turning on the light), sensing only when a user pushes a physical button on the device, or they do not know (see Appendix B-C).

9) *Demographic Questions*: We asked general demographic questions to capture participants’ age, gender, highest degree earned, and background in technology (see Appendix B-D).

B. Data Analysis

Our study used a repeated-measures design in which participants were presented with multiple scenarios with the same type of questions. This design results in multiple observations for each participant that are not completely independent. Therefore, we used a statistical method that allows us to use random effects to account for these dependencies. To quantitatively measure the significance of our independent variables, we used Cumulative Link Mixed Models (CLMMs) with logit as the link function. This allowed us to model all five levels of our ordinal Likert scale responses for our dependent variables (risk perception and willingness to purchase) rather than having to

bin them into two levels as required by logistic regression [68]. We used a significance threshold of 0.05.

We used content analysis [69] to find the reasons participants’ risk perception and willingness to purchase were impacted or not impacted by privacy and security information. The first author constructed a codebook used to analyze free-text responses. The first author and another researcher independently applied the codebook to the responses and through several meetings, iteratively revised the codebook. After discussing the coding discrepancies and resolving major disagreements, we reached a Cohen’s Kappa inter-coder agreement of 81%, which is considered *excellent* [70]. In case of unresolved disagreements, we report on the findings of the first author.

C. Limitations

We tested the impact of privacy and security attributes on participants’ *self-reported* risk perception and willingness to purchase. While these measures have been shown to strongly correlate with actual behavior [27]–[30], [71], they are not a complete substitute for a study that observes real purchase scenarios. We expect that our approach likely exaggerates the effect that each attribute-value has on risk perception and willingness to purchase, as these attribute-values will compete with each other for a consumer’s attention in a real purchase scenario. Thus, in a more realistic study we would expect attribute-value pairs that exhibited a minor effect in our study to exhibit little or no effect.

We designed our study so that we could measure the effectiveness of each privacy and security attribute-value pair in isolation. This allowed us to study the impact of each attribute in order to prioritize the information that is included on an IoT label as well as to identify misconceptions associated with individual attributes. However, a full privacy and security label would include more than one attribute. Further testing is needed to explore the nuances in consumers’ risk perception and willingness to purchase when presented with a complete IoT privacy and security label. Again, we expect that the effect of each individual attribute will be muted when presented in the context of a complete label. However, interaction effects may also emerge.

We evaluated the importance of a limited number of factors in describing risk perception and willingness to purchase. For instance, we tested only two types of IoT devices, three types of recipients, and two extreme values for the tested security and privacy attributes. It would be useful to also test other levels of these factors, for example, gifting a device to a child who is being protected by Children’s Online Privacy Protection Act (COPPA). It would also be useful to test the common values of privacy and security attributes that fall in between the extremes, e.g., default passwords and user-changeable passwords rather than just no passwords and multi-factor authentication.

Our survey questions might have primed participants to think about privacy, security, and risks more than they would in realistic purchase scenarios. However, as participants were equally primed in all conditions, we expect little impact on

the *relative* risk and desire to purchase for each attribute-value and the relative ordering between them.

Our methodology and regression analysis capture the impact of each attribute-value pair in isolation. However, as multiple attribute-value pairs were presented to each participant, interaction effects among these factors may exist that we did not investigate.

We provided a plain-language consumer explanation for each privacy and security attribute-value pair to help participants understand what they mean. To assess participants' attention to the provided explanations, we asked an attention-check question and asked participants about their level of confidence in understanding the attribute-values. However, we did not test participants on their knowledge related to the privacy and security implication of each attribute-value pair and thus have no quantitative assessment of whether participants' level of confidence in their knowledge correlates with their actual understanding. However, our qualitative analysis of open-ended responses indicated that participants did not seem to find the attribute-value pairs confusing. Even so, as we will discuss in Section IV-E, we found some misconceptions about the implications of a few privacy and security attribute-value pairs. More detailed knowledge questions are needed to fully assess participants' understanding of each attribute-value pair.

Finally, we recruited participants from MTurk who reside in the US. Residents of other countries likely have different perspectives on privacy and security risk and willingness to purchase IoT devices. Furthermore, our participants are not completely representative of the US population. For instance, as shown in Table IV in Appendix D, our participants were slightly more educated and younger than the general US population. Past studies have found that, even when controlling for these demographic factors, MTurk workers may have heightened privacy concerns compared with the US population [72].

IV. RESULTS

We first describe our participants and present summary statistics. Next we present our statistical models of risk perception and willingness to purchase. Finally, we provide insights from our qualitative analysis of participants' responses.

A. Participants and Study Conditions

We initially recruited 1,710 MTurk participants and excluded those whose answers for all our open-ended questions were irrelevant. This resulted in 1,371 participants who are included in our analysis. All of these participants answered at least two out of their three attention-check questions correctly. Overall, at least 90% of participants correctly answered the attention-check questions for all but two of the 33 attribute-value pairs, indicating that participants were paying attention to the label information we presented to them. The two attribute-value pairs with the most wrong answers were: *security audit: internal & external* (22% incorrect), and *control over: device storage* (21% incorrect).

We randomly assigned each participant to one out of six study conditions (based on two device types and three device

recipients). Each participant was asked to answer questions related to three random attribute-value pairs for the condition they were assigned to. There were between 224 and 233 participants per condition and between 119 and 132 participants per attribute-value pair.

Our participants were 54% male and 45.5% female. Compared to the 2018 US Census data [73], participants were younger and better educated. Participant demographic information is provided in Appendix D.

B. Summary Statistics

1) *Concern Level*: We found a strong correlation between the device type (smart speaker, smart light bulb) and the level of concern with the type of device (binary variable¹), $\chi^2(1, N = 1371) = 189.14, p < 0.001, \phi = 0.37$ [74]. 62% of the participants who were assigned to the smart light bulb conditions reported being concerned, mainly due to the unforeseeable consequences of their data being accessed by unauthorized parties. In the smart speaker conditions, 93% reported being concerned about these devices. Most participants mentioned that they are concerned about smart speakers always listening to them. The difference in participants' level of concern for the smart speaker and smart light bulb is consistent with our hypothesis, as well as past findings [5], [26], [61].

2) *Purchase History*: 54% of participants reported having a smart speaker in their home, and among those, 53% had purchased the device themselves. Only 12% of participants reported having a smart light bulb, and 61% of those reported that they had purchased it themselves.

Among those who did not have the smart device in question, 23% reported that they had been in the market to purchase it earlier. The main reasons stated for not going through with the purchase were their price (30% for the smart speaker and 48% for the smart light bulb) and lack of necessity (44% and 34%, respectively). Privacy and security concerns were also reported by 26% and 9% of participants as reasons not to purchase the smart speaker and the smart light bulb, respectively.

3) *Confidence Level in Understanding Label Information*: More than 70% of participants reported being somewhat, moderately, or very confident about knowing what the label information meant for all but two attributes. Participants' level of confidence was significantly lower (p -value < 0.05) for security audit and data linkage.²

C. Device Functionality Perception

We found that most participants have a correct understanding about how the smart device in their study condition works. In the smart light bulb condition, 72% of participants believed that the light bulb always senses whether someone is present in the room, 12% reported not being sure how the device works, 10% believed that the device starts sensing when a button on the device is pushed, and 6% thought that the smart light bulb

¹We coded "not at all concerned" as 0 and "only slightly concerned," "somewhat concerned," "moderately concerned," and "very concerned" as 1.

²We constructed a CLMM to model the impact of attribute-value on the level of confidence.

starts sensing when triggered by someone’s presence in the room. In the smart speaker condition, we found that 53% of participants had a belief that the device waits for the user to mention the wake word (e.g., “Alexa”, “OK Google”), 39% thought that the device is always sensing, 4% reported not knowing how the smart speaker works, and 4% believed that the device starts listening when the user presses a button to turn on the device microphone.

D. Risk Perception and Willingness-to-Purchase Models

We were interested in understanding the impact of various factors on two dependent variables (DVs): participants’ risk perception and willingness to purchase the smart device. We built two Cumulative Link Mixed Models (CLMMs) to describe our DVs. The factors we included in each model are as follows:³

- `sp_attribute_value`: 33 security/privacy attribute-value pairs (see Table I). Among these 33 attribute-value pairs, only three of them were randomly selected and presented to each participant, while the rest of the attributes were *absent*, i.e., not shown to the participant.
- `prior_scenarios`: Number of prior scenarios seen by that participant, with three levels: 0, 1, and 2 scenarios.
- `device_exposure`: How much exposure participants have to the smart device, with three levels: “Not having the device,” “purchased the device” (owned the device and purchased it themselves), and “didn’t purchase the device” (owned the device, but did not purchase it themselves).
- `device_type`: Type of the device, with two levels: Smart speaker and smart light bulb.
- `device_recipient`: Who the device is being purchased for, with three levels: Yourself, friend, and family.

Our dataset included three scenarios from each of the 1,371 participants for a total of $N = 4,113$ observations. We asked participants to specify, on a Likert scale, the impact of each presented attribute-value pair on risk perception and willingness to purchase (see Appendices B-B1 and B-B2), leading to $J = 5$ ordinal response categories. We modeled risk perception and willingness to purchase by fitting two CLMMs to the data. In these CLMMs, we included a random intercept per participant to take the natural and inherent variations among different participants into account. In each model, the probability that the i^{th} observation, $i \in \{1, \dots, N\}$, falls in the j^{th} response category or below, $j \in \{1, \dots, J-1\}$, is modeled as

$$\begin{aligned} \text{logit}(\Pr(Y_i \leq j)) = & \alpha_{j|j+1} - u_{\text{participant}_i} \\ & - \beta_{\text{sp_attribute_value}_i} \\ & - \beta_{\text{prior_scenarios}_i} - \beta_{\text{device_exposure}_i} \\ & - \beta_{\text{device_type}_i} - \beta_{\text{device_recipient}_i}, \end{aligned} \quad (1)$$

³We also initially included the demographic information (see Appendix D) and the type of prior privacy and security attribute-value, but we found their impact to be insignificant in both models. Therefore, we decided to remove them from our final models to better fit the data (according to Akaike Information Criterion (AIC), which we used to assess the model goodness of fit [75]). Except for demographic factors as well as the prior privacy and security attribute-value, removing other factors from the models resulted in a decline in model fit. Therefore, we did not remove any of the remaining factors from our models.

where $\alpha_{j|j+1}$ denotes the threshold parameter between response categories j and $j+1$, and $u_{\text{participant}_i}$ denotes the random effect for the participant in the i^{th} observation, modeled as an independent and identically distributed (i.i.d.) Gaussian random variable with zero mean and variance of σ_u^2 , i.e., $u_{\text{participant}_i} \sim \mathcal{N}(0, \sigma_u^2), \forall i \in \{1, \dots, N\}$. Moreover, for each factor in the model, β_{factor_i} denotes the model coefficient corresponding to the level of that factor present in the i^{th} observation. Note that in the case of security/privacy attribute-value pairs, our model captures i) whether each of the 16 security and privacy attributes in Table I was present or absent in each observation, and ii) if present, whether the most protective or the least protective value of that attribute was observed.⁴

Table II presents the results of the two models. For each model, we present the variance of random effects σ_u^2 and threshold parameters $\{\alpha_{j|j+1}\}_{j=1}^{J-1}$. Moreover, for each factor, we report its estimate, i.e., the corresponding β coefficient in (1), as well as its standard error and p -value. We also provide the *odds ratios of increased risk* and *willingness to purchase* for each factor, respectively defined as

$$\text{OR}_{\text{risk}(+)|\text{factor}} \triangleq \frac{\Pr(\text{slightly/strongly increased risk} | \text{factor})}{1 - \Pr(\text{slightly/strongly increased risk} | \text{factor})} \cdot \frac{\Pr(\text{slightly/strongly increased risk} | \text{factor_baseline})}{1 - \Pr(\text{slightly/strongly increased risk} | \text{factor_baseline})}, \quad (2)$$

$$\text{OR}_{\text{purchase}(+)|\text{factor}} \triangleq \frac{\Pr(\text{slightly/strongly increased willingness to purchase} | \text{factor})}{1 - \Pr(\text{slightly/strongly increased willingness to purchase} | \text{factor})} \cdot \frac{\Pr(\text{slightly/strongly increased willingness to purchase} | \text{factor_baseline})}{1 - \Pr(\text{slightly/strongly increased willingness to purchase} | \text{factor_baseline})}. \quad (3)$$

Finally, we provide the *odds ratios of decreased risk* and *willingness to purchase* for each factor, respectively defined as

$$\text{OR}_{\text{risk}(-)|\text{factor}} \triangleq \frac{1}{\text{OR}_{\text{risk}(+)|\text{factor}}}, \quad (4)$$

$$\text{OR}_{\text{purchase}(-)|\text{factor}} \triangleq \frac{1}{\text{OR}_{\text{purchase}(+)|\text{factor}}}. \quad (5)$$

In both models, we select *purpose: device function* to be the baseline for `sp_attribute_value`, as it is the purpose that most IoT devices will serve, possibly in addition to others. The *smart light bulb* is selected as the baseline for `device_type`, since its data collection is less concerning, the baselines for factors `device_exposure` and `device_recipient` are selected to be the most common values of these factors, and the baseline for `prior_scenarios` is *0 scenarios* as its first level. The selection of baselines implies that their corresponding β coefficients in (1) are set to zero by the model. Note that the selection of baselines will not affect the final output of the models, e.g., in terms of the cumulative response category probabilities (1) and odds ratios (2)-(5).

For the risk perception model, a positive estimate for a factor indicates an increase in risk perception compared to the baseline of that factor. Similarly, in the willingness to purchase

⁴In case of the presence of the *control over* attribute, our model captures which of three most protective values was observed by the participant.

Row	Factor	Risk perception					Willingness to purchase					
		OR ₍₊₎	OR ₍₋₎	Estimate	Std. Error	p-value	OR ₍₊₎	OR ₍₋₎	Estimate	Std. Error	p-value	
Least protective	sp_attribute_value (baseline = purpose: device function)											
	1	Sold to: third parties	26.58	0.04	3.28	0.34	***	0.04	24.05	-3.18	0.30	***
	2	Access control: none	19.11	0.05	2.95	0.31	***	0.06	16.28	-2.79	0.28	***
	3	Shared with: third parties	13.07	0.08	2.57	0.29	***	0.09	10.80	-2.38	0.26	***
	4	Sharing frequency: continuous	10.59	0.09	2.36	0.29	***	0.16	6.23	-1.83	0.25	***
	5	Cloud retention: indefinite	10.38	0.10	2.34	0.28	***	0.12	8.50	-2.14	0.26	***
	6	Device retention: indefinite	9.97	0.10	2.30	0.28	***	0.16	6.30	-1.84	0.26	***
	7	Security update: none	9.58	0.10	2.26	0.29	***	0.12	8.41	-2.13	0.27	***
	8	Device storage: identified	8.84	0.11	2.18	0.28	***	0.15	6.69	-1.90	0.26	***
	9	Cloud storage: identified	6.42	0.16	1.86	0.27	***	0.18	5.42	-1.69	0.25	***
	10	Average time to patch: 6 months	6.36	0.16	1.85	0.26	***	0.21	4.76	-1.56	0.25	***
	11	Purpose: monetization	6.11	0.16	1.81	0.26	***	0.11	8.94	-2.19	0.26	***
	12	Data linkage: internal & external	5.70	0.18	1.74	0.27	***	0.19	5.16	-1.64	0.25	***
	13	Inference: additional info	5.58	0.18	1.72	0.28	***	0.17	5.87	-1.77	0.26	***
	14	Security audit: none	5.26	0.19	1.66	0.30	***	0.24	4.14	-1.42	0.25	***
15	Collection frequency: continuous	4.81	0.21	1.57	0.26	***	0.23	4.35	-1.47	0.25	***	
Most protective	16	Average time to patch: 1 month	3.13	0.32	1.14	0.26	***	0.34	2.97	-1.09	0.25	***
	17	Security audit: internal & external	0.37	2.72	-1.00	0.27	***	1.62	0.62	0.48	0.25	*
	18	Device storage: none	0.11	8.76	-2.17	0.27	***	5.99	0.17	1.79	0.25	***
	19	Inference: none	0.11	9.49	-2.25	0.26	***	3.90	0.26	1.36	0.24	***
	20	Security update: automatic	0.10	9.87	-2.29	0.26	***	4.26	0.23	1.45	0.24	***
	21	Sharing frequency: on user demand	0.08	13.07	-2.57	0.26	***	6.23	0.16	1.83	0.24	***
	22	Control over: cloud data deletion	0.07	14.15	-2.65	0.26	***	5.99	0.17	1.79	0.25	***
	23	Collection frequency: on user demand	0.07	14.01	-2.64	0.25	***	7.03	0.14	1.95	0.24	***
	24	Control over: device retention	0.07	14.01	-2.64	0.25	***	7.24	0.14	1.98	0.24	***
	25	Data linkage: none	0.07	15.33	-2.73	0.26	***	6.23	0.16	1.83	0.24	***
	26	Sold to: none	0.05	20.70	-3.03	0.26	***	11.13	0.09	2.41	0.25	***
	27	Cloud storage: none	0.05	22.20	-3.10	0.26	***	6.62	0.15	1.89	0.25	***
	28	Control over: device storage	0.04	28.50	-3.35	0.26	***	11.25	0.09	2.42	0.25	***
	29	Device retention: none	0.02	41.26	-3.72	0.27	***	13.20	0.08	2.58	0.25	***
	30	Access control: MFA	0.02	45.60	-3.82	0.27	***	8.25	0.12	2.11	0.24	***
	31	Shared with: none	0.02	50.91	-3.93	0.27	***	18.54	0.05	2.92	0.25	***
	32	Cloud retention: none	0.02	50.91	-3.93	0.27	***	13.74	0.07	2.62	0.25	***
device_type (baseline = smart light bulb)												
33	Smart speaker	1.55	0.64	0.44	0.08	***	0.66	1.52	-0.42	0.07	***	
device_recipient (baseline = yourself)												
34	Friend	1.01	0.99	0.01	0.08	0.83	0.88	1.14	-0.13	0.08	*	
35	Family	0.97	1.03	-0.03	0.08	0.69	1.00	1.00	0.00	0.08	0.97	
device_exposure (baseline = not having the device)												
36	Didn't purchase the device	0.79	1.27	-0.24	0.11	*	1.68	0.59	0.52	0.10	***	
37	Purchased the device	0.63	1.60	-0.47	0.10	***	2.03	0.49	0.71	0.10	***	
prior_scenarios (baseline = 0 scenarios)												
38	1 scenario	1.19	0.84	0.17	0.08	*	0.76	1.32	-0.28	0.07	***	
39	2 scenarios	1.14	0.88	0.13	0.08	0.10	0.63	1.58	-0.46	0.07	***	
threshold coefficients												
40	1 2	-	-	-3.53	0.21	-	-	-	-1.83	0.19	-	
41	2 3	-	-	-1.59	0.20	-	-	-	-0.80	0.19	-	
42	3 4	-	-	0.08	0.19	-	-	-	1.60	0.19	-	
43	4 5	-	-	1.40	0.20	-	-	-	3.41	0.20	-	
random effects												
44	σ_u^2	-	-	0.27	-	-	-	-	0.43	-	-	
Note: * $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$												

TABLE II: We used CLMM and built two models to identify the significance of various factors in changing participants' risk perception and willingness to purchase. The Nagelkerke R^2 values for the risk perception and willingness to purchase models are 0.74 and 0.68, respectively. The Cox & Snell values for the risk perception and willingness to purchase models are 0.71 and 0.65, respectively [76]. For the security and privacy attribute-value pairs, except for the *control over* attribute, our models capture three levels of each attribute, i.e., most protective, least protective, and absent (i.e., not shown), while they capture four levels for the *control over* attribute, namely its three most protective levels as well as absent.

model, a positive estimate indicates an increase in participants' desire to purchase the smart device, and a negative estimate indicates hesitance to purchase, all compared to the baseline.

Since we showed three scenarios to each participant, there might exist two-way or three-way interaction effects among the presented attribute-value pairs. An interaction term is statistically defined between levels of multiple distinct factors, while in our constructed model, attribute-value pairs are the levels of the same factor (*sp_attribute_value*). Therefore, our model is not able to capture such potential interactions. As full privacy and security labels would include

multiple attribute-value pairs, future studies should carefully explore the interactions among the presented factors.

Privacy and Security Information. In both models, all the privacy and security attribute-value pairs significantly changed participants' risk perception and willingness to purchase. For almost all of these pairs, the direction of the change was aligned with our hypothesis (see Table I), except for the average time to patch. The Underwriters Lab (UL) guidelines suggest that the most severe vulnerabilities should be patched within 1 month and less severe vulnerabilities within 3 months [77]. Thus, we hypothesized that participants' perceived risk would decrease

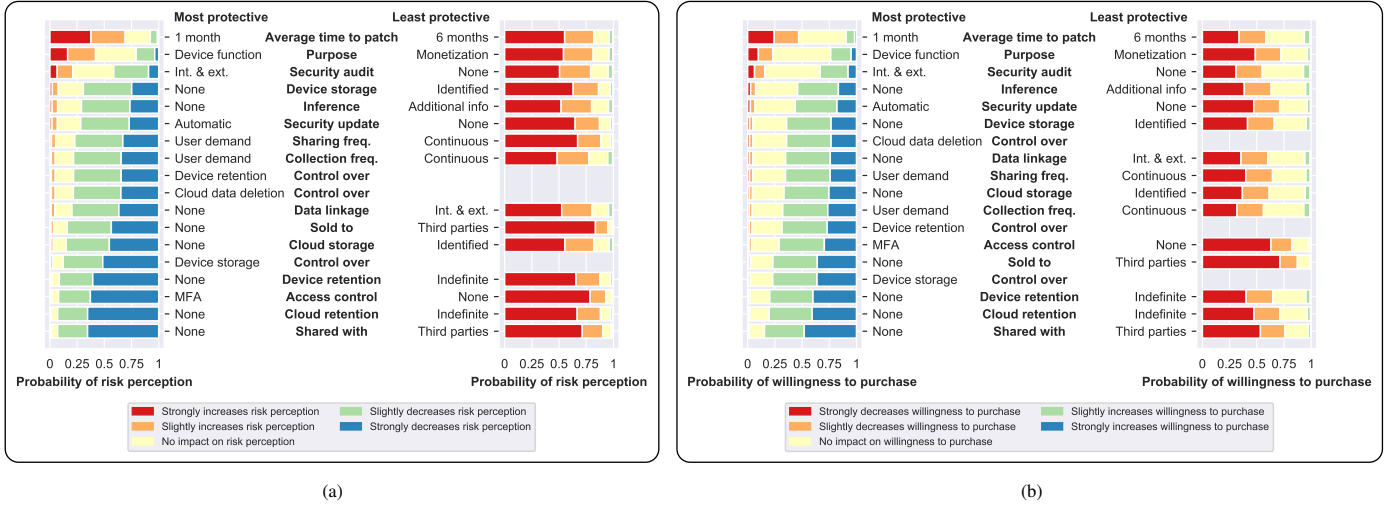


Fig. 1: Based on the CLMM parameters in the baseline condition (i.e., all the model factors, except `sp_attribute_value`, being at their baseline values), we computed and plotted the probabilities of each privacy and security attribute-value pair increasing, decreasing, or having no impact on risk perception (Left Fig: 1a) and willingness to purchase (Right Fig: 1b).

if vulnerabilities were patched within 1 month, and increase if they were patched within 6 months. However, the findings from the CLMMs show that average time to patch of both 1 month and 6 months strongly increase the perceived risk and decrease willingness to purchase (see Table II, rows 16 and 10). In fact, most of our participants reported that they would like the company to patch vulnerabilities within a few days.

The CLMM results show that data being sold to third parties (Table II, row 1) and having no control over access (Table II, row 2) had the biggest impact on increasing participants' risk perception, whereas no cloud retention (Table II, row 32) and data not being shared with third parties (Table II, row 31) had the biggest impact on decreasing the perceived risk. The direction of the impact of attribute-value pairs in the willingness to purchase model was similar to the risk perception model. However, the relative importance of the attribute-value pairs were not exactly the same across the two models. For instance, knowing information would not be shared most influenced willingness to purchase while no cloud retention most influenced risk perception.

Comparing the odds ratios in Table II, we observe that for all the least protective values, the odds ratios of increasing risk perception are higher than their corresponding odds ratios of decreasing the desire to purchase. Similarly, for all the most protective attribute-value pairs except *average time to patch: 1 month*, the odds ratios of decreasing risk perception are higher than their corresponding odds ratios of increasing willingness to purchase. This implies that the tested privacy and security attribute-value pairs were more powerful in changing participants' risk perception than in changing their willingness to purchase. From the open-ended responses, we observe a possible explanation: Participants report that privacy and security are among many factors they consider when purchasing a device. Several participants who reported that a privacy and security attribute-value does not have an impact on their willingness to purchase mentioned that their willingness to

purchase the smart device is driven by its price or features.

Figure 1 illustrates the probabilities of the five response categories for risk perception and willingness to purchase based on the CLMM estimates. Our analysis shows that participants are significantly more likely to specify an increase rather than a decrease in risk perception for all the least-protective values. The reverse trend was also mostly true for the most protective values. There were, however, a few exceptions. We found that the attribute-value pairs *security audit: internal & external*, *purpose: device function*, and *average time to patch: 1 month* led to a considerable probability of increased risk perception (21%, 41%, and 68%, respectively), suggesting that these attribute-value pairs may not be clear. We discuss the open-ended responses that provide insights into these unexpected findings in Section IV-E. As the figure shows, unlike their most protective values, the least protective values *purpose: monetization* and *average time to patch: 6 months* had a large impact on increasing risk perception and decreasing willingness to purchase.

Figure 2 shows jitter (scatter) plots of participants' perceived risk levels and willingness to purchase when presented with attributes alongside their most and least protective values. As the plots demonstrate, the correlation between risk perception and willingness to purchase differs based on the attribute. For instance, Figure 2a shows that most participants perceived multi-factor authentication (MFA) as decreasing risk (89%) and no access control as increasing risk (97%). While this was generally consistent with their willingness to purchase, the figure shows that some participants who perceived MFA as risk reducing were actually no more likely or even less likely to purchase a device with MFA (31%). Our qualitative responses suggest this is mainly due to MFA usability challenges. Likewise, Figure 2b shows that most participants perceived no sharing as decreasing risk (85%) and sharing with third parties as increasing risk (95%). However, in this case, risk perception was much more likely to be correlated with willingness to

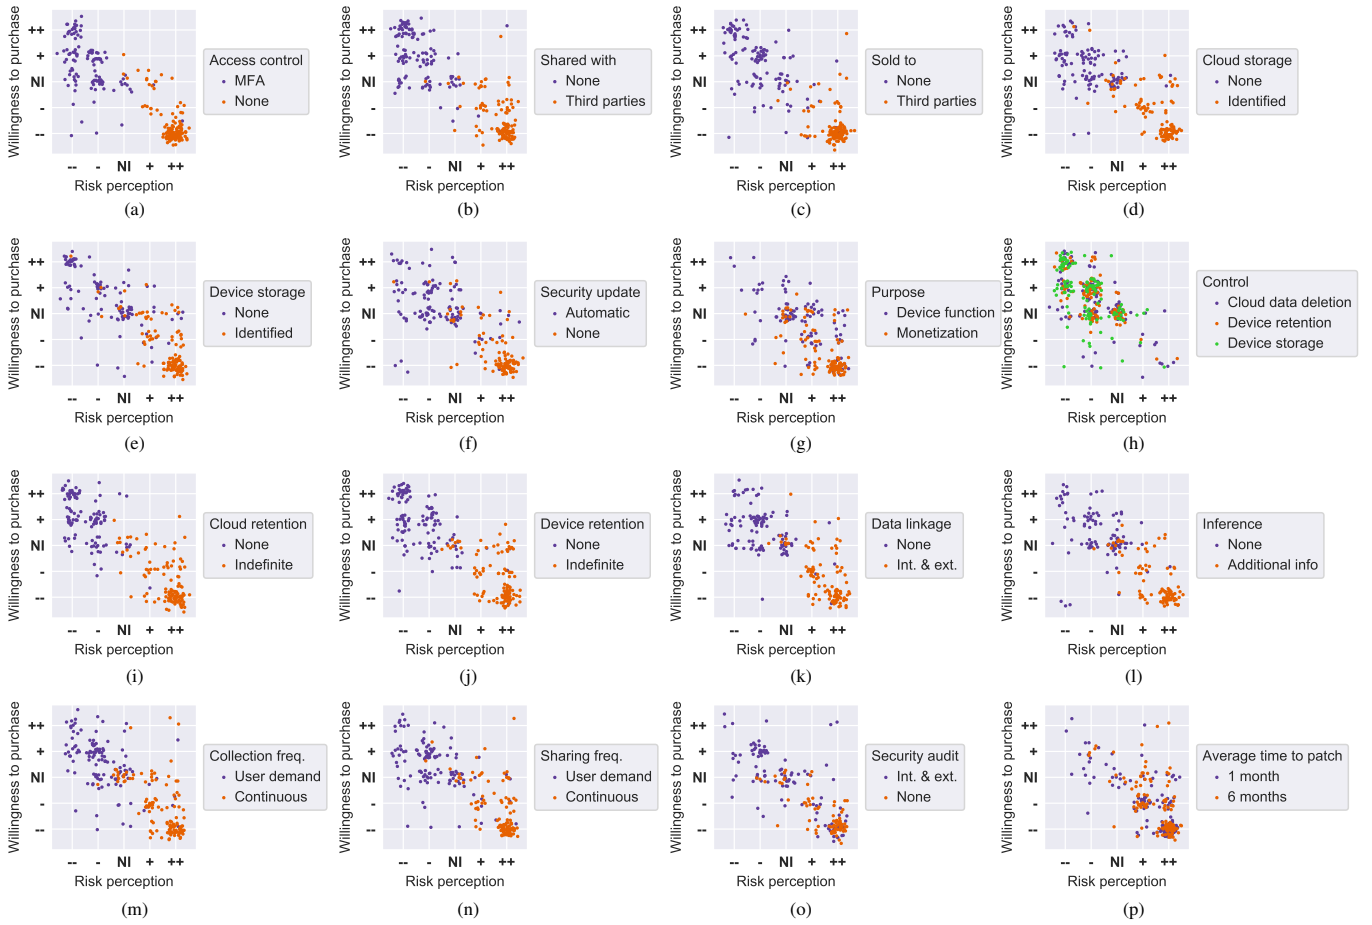


Fig. 2: Jitter (scatter) plot of participants’ willingness to purchase vs. perceived risk for all 16 privacy and security factors and levels. We use the following notation to label the x and y axes: -- represents “strongly decrease,” - represents “slightly decrease,” NI represents “no impact,” + represents “slightly increase,” and ++ represents “strongly increase.”

purchase. On the other hand, Figure 2p shows that participants perceived both values of average time to patch as risky (78% for 1 month and 91% for 6 months) and also decreasing their likelihood to purchase (67% for 1 month and 75% for 6 months). Finally, as we hypothesized, Figure 2h confirms that all levels of control seem to reduce participants’ perceived risk and increase their willingness to purchase.

Type of the Device. In addition to the privacy and security attributes, our results show that the `device_type` was a significant factor to describe risk perception and willingness to purchase. In particular, compared to a smart light bulb (baseline for `device_type`), participants perceived a significantly higher risk for a smart speaker (estimate = 0.44, p -value < 0.001) and they were significantly less willing to purchase the smart speaker (estimate = -0.42, p -value < 0.001).

Recipient of the Device. The `device_recipient` was not a statistically significant factor to describe risk perception. However, participants’ willingness to purchase the device significantly decreased when the recipient of the device was a friend (estimate = -0.13, p -value < 0.05) compared to purchasing the device for themselves (baseline for `device_recipient`). The most common reason participants mentioned for this

decrease was that participants thought friends should decide for themselves if they want to use the device based on the privacy and security information. However, in the study conditions where we asked participants about family members as the device recipients, they reported feeling more responsible for purchasing a safe device.

Device Exposure. In each study condition, we asked participants to specify whether they owned a device of the type in that condition and how they acquired it (see Appendix B-A). Similar to prior work [78]–[81], our analysis shows that compared to not having the device, participants who had the device, either by purchasing the device themselves (estimate = -0.47, p -value < 0.001) or acquiring it in other ways (estimate = -0.24, p -value < 0.05), perceived a significantly lower risk. Moreover, our results indicated that those participants who had the device had a significantly higher (p -value < 0.001) desire to purchase a future device, compared to those who did not have the device. Nevertheless, we found no significant differences based on how participants acquired the device.

Prior Scenarios. Our CLMMs show that the number of prior scenarios influence the risk perception and willingness to purchase for the current scenario. More specifically,

we found that compared to the first scenario (baseline for *prior_scenarios*), participants perceived a significantly higher risk for the second scenario (estimate = 0.17, p -value < 0.05). In the willingness to purchase model, participants expressed a significantly lower desire to purchase the smart device in the second (estimate = -0.28, p -value < 0.001) as well as the third scenarios (estimate = -0.46, p -value < 0.001) compared to the first scenario. We found a significant difference in willingness to purchase between the second scenario and the third scenario (p -value < 0.05). One possible explanation for this effect could be that participants became more sensitive to risks as they were exposed to more scenarios and, therefore, gradually became less willing to purchase the device.

E. Qualitative Results

In most cases, open-ended responses echo and provide additional insights into our quantitative results. For example, the open-ended responses help explain why participants were so concerned about access to their information or devices.

Participants frequently expressed concerns about not knowing how their information would be used if it was shared with or sold to third parties. P8 reported: "Once data is sold, you do not know where or what it is being used for, or if it will be sold again. This will enable tracking by companies you are unaware of." For similar reasons, participants were more inclined to purchase devices that did not share their data or store their data in the cloud.

Participants were also concerned about unauthorized access to devices that lacked access controls. P280 said: "Having no control means that I have no way of keeping people from accessing my light bulb and the sensor. There is no password protection or authentication methods where it would allow only myself to access the software."

The open-ended explanations also helped us identify a number of common reasons for responses that differed from what we had hypothesized, explained below.

Insufficient Information. Not having enough information was the most common reason that five of the factors did not decrease some participants' risk perception, despite our hypothesis. In their open-ended responses, participants requested additional information:

- *Average time to patch: 1 month:* Out of 122 participants, 68 reported being unsure about whether one month is too short or too long, expressing that they need more information on why it takes manufacturers one month to fix vulnerabilities.
- *Security audit: internal & external:* Out of 125 participants, 31 wanted to know who the auditors are and what information they have access to when conducting audits.
- *Data linkage: internal & external:* Out of 121 participants, 22 wanted to know what these internal and external data sources are and what the consequences of their data being linked with such data sources would be.
- *Security update: automatic:* Out of 128 participants, 21 wanted to know how often their device would get updated and whether they can control this frequency.

- *Device storage: none:* Out of 123 participants, 17 wanted to know whether information would be stored on the cloud or if no data would be stored at all.

Lack of Trust in Manufacturers. The second most common reason that a factor did not impact participants' risk perception and willingness to purchase was not trusting manufacturers to follow their stated practices. Some participants (53/123) expressed a lack of trust when they saw *purpose:device function*, for which we provided the explanation "Data is being collected to provide the main device features, improve services, and help develop new features." Although we hypothesized that providing data for device functionality should decrease the perceived risk, that was true for only 12% of participants seeing that information (15/123). Other participants stated that this information would not impact their risk perception or would even increase risk, mostly due to their lack of trust in manufacturers. As P778 explained: "The companies who collect data are incredibly untrustworthy. They do not have consumers' best interests in mind when they are utilizing the data they collect."

On the other hand, lack of trust also resulted in most participants (101/124) expressing concern about *purpose: monetization*. P298, who perceived a strongly higher risk, reported: "This means that the device is collecting all kinds of information. The device is sharing all kind of information with companies you do not know if they can be trusted."

A few participants (11/128) mentioned lack of trust when assessing the risk of *security update: automatic*. They reported that manufacturers can apply unwanted changes to their devices under the pretext of security updates.

Shared with: none and *sold to: none* were other attribute-value pairs where participants expressed lack of trust. Some participants mentioned that while they believe this would decrease potential privacy risks, they do not trust the manufacturers not to send their data to other companies for profit.

Participants' comments about trust are consistent with prior work that has identified the trust people have in organizations as one of the factors affecting risk perception [82]–[87].

Following Standard Practices. For eight attribute-value pairs, participants believed that the reported privacy and security practices were standard, and, therefore, having them would not provide additional protection. This reason was most commonly mentioned about the attribute-value pairs that conveyed to participants a sense of control over their devices.

These pairs were *control over: cloud data deletion* (27/122), *control over: device retention* (21/125), *control over: device storage* (19/128), *security audit: internal & external* (14/125), *collection frequency: on user demand* (12/132), *sharing frequency: on user demand* (10/127), *security update: automatic* (9/128), and *access control: multi-factor authentication* (5/127).

P878 believed that data collected with user consent is standard: "I would assume this is standard and normal. If the company is not ethical, they will just collect the data anyway." Moreover, P205 reported that knowing about internal and external security audits will not change their willingness to purchase the device by saying: "I'm pretty sure audits are

standard. I assume that any product reaching consumer shelves has been audited by at least someone.”

Usability Challenges. Some participants reported that requiring users to use *multi-factor authentication* (37/127) or consent to data collection each time *on user demand* (27/132) would affect device usability. P1334 was particularly concerned about MFA for shared in-home devices: “Accessing the device via authentication would then become a hassle and/or annoying. For instance, what if my wife or a guest wanted to use the speaker?”

All participants who mentioned usability challenges reported that the corresponding attribute-value pairs would decrease the risks. However, they reported that those attribute-value pairs would have no impact on, and in most cases would even decrease, their willingness to purchase.

Desire to have Control. Automatic security updates are widely recommended [63], [64], [77], [88] and, as expected, *security update: automatic* significantly decreased participants’ risk perception. However, for some participants (19/128), this decrease in risk perception did not positively impact their willingness to purchase due to the lack of control implied by this practice. P535 reported: “I want to have full control over updating my devices to decrease the risk of installing an update that has a security flaw.”

No Initial Concerns. In the study conditions where the device was a smart light bulb, many participants exhibited no initial concern (those who answered “not at all concerned”). 57% of those participants said they were not aware of any consequences of the data collection (129/255). Therefore, privacy and security attribute-values that generally reduced risk perception had no impact on some participants because they did not perceive a risk to begin with.

We found a strong correlation between type of device and the concern level, and as we previously mentioned, *device_type* was a significant factor in explaining participants’ risk perception (see Table II).

Misconceptions. Although we did not find any indication of confusion over the presented attribute-value definitions, there were two attribute-value pairs that a few participants appeared to have misconceptions about, thus affecting their responses. One pair was *security update: none*. A few participants (5/125) mentioned that receiving no security updates implied maximum security protection, alleviating the need for updates. Another misunderstood attribute was the average time to patch. Some participants (8/253 participants across both values of average time to patch) mentioned that a device that receives security patches must not be secure or it would not need patches. For instance, P906 mentioned: “On the label, it advertises that patches are even needed. That is why there is a perception of decreased privacy.”

Other Reasons. There were some reasons that were mentioned by only a small number of participants. One of the reasons that an attribute-value pair did not change participants’ willingness to purchase was that they had already decided either to purchase or not to purchase that kind of device (68/1371 across all study conditions) due to factors such as

its functionality or their prior privacy and security concerns with the device. P750, who was asked to imagine purchasing a smart speaker, reported: “There is little incentive for the companies to keep data secure. The fact that IoT devices all send data to a privately controlled server is unacceptable. Any low-level employee or barely motivated hacker could get access to all the information. The government could just ask for the information. I don’t want any such devices in my home.”

Some participants who were in the study conditions, where they were asked to imagine purchasing the smart device for a family member or a friend as a gift, mentioned that they did not feel comfortable making a decision for the gift recipient without knowing their preferences (49/915 participants across conditions with friends or family as *device_recipient*).

Finally, some participants (82/1371 across all study conditions) stated that as long as data is being collected by the device, no privacy and security practice could eliminate the potential risks. For instance, P1338 reported: “Just because I can control how the data is retained on the device doesn’t mean I have any control about how that data is collected and how it is used while it is retained - the company could still upload the data from the device to their server before it is deleted from the device.”

V. DISCUSSION

Our findings pave the path to an improved IoT privacy and security label, which can ultimately lead to a safer and more secure IoT ecosystem. In this section, we start by providing recommendations on how to further enhance information communication and effectiveness of IoT labels. We then provide a discussion on the impact of risk acceptability on consumers’ willingness to purchase. We continue with a discussion of how labels may impact consumers’ purchase behavior. Finally, we provide a brief discussion on the future of IoT labels and suggest paths toward adoption.

A. Toward a More Effective IoT Label Design

Our quantitative and qualitative results suggest that presenting privacy and security attributes in an IoT label is likely to influence participants’ risk perception and willingness to purchase. Although almost all of the tested attribute-value pairs were statistically significant in explaining risk perception and willingness to purchase, based on qualitative responses we propose several ways to better convey risk to consumers and help inform their purchase behavior. Our proposed changes should be tested in future user studies.

1) *Reducing Information Uncertainty:* Our qualitative findings showed that understanding the provided definitions of attribute-values (see Appendix C) is not always enough to make an informed decision. In Section IV-E, we mentioned several additional pieces of information that participants thought could help inform their decisions.

Based on these qualitative responses, we recommend that manufacturers provide the following additional information:

- Justification as to why the manufacturer has a specific privacy and security practice in place.

- In what ways a specific privacy and security practice could protect or harm consumers.
- What controls consumers have related to each privacy and security attribute.
- If an option is being offered to control a specific privacy and security practice, what steps users need to take to enable that option.

It is important to note that although adding information to the label would mitigate consumers' uncertainty, too much information could overwhelm consumers. Therefore, we propose adding extra information on the label in an expanded view accessed by a plus sign that is placed next to each attribute. Note that the expanded view is only available on the secondary layer of the label, which is in an online-only format.

2) *Placement of Information on the Label*: Our quantitative analysis suggests that the proposed distribution of attributes between the primary and secondary layers of the label [13] is mostly appropriate based on each attribute's impact on risk perception and willingness to purchase. We found that all the primary-layer attributes we tested accurately communicated increased or decreased risk to participants, with the exception of *purpose: device function* (see Figure 1). As discussed in Section IV-E, participants expressed a lack of trust in manufacturers to use collected data only to enable the main device features. However, we expect that manufacturers could increase trust by being more explicit about the purposes for which they use data. This explanation could be added to the secondary layer of the label. In addition, future work should explore whether this information is useful enough to consumers that it is worth adding to the primary layer. For some types of sensors/devices, this information will be fairly obvious (e.g., motion detector sensing motion), but for others it may be more surprising (e.g., motion detector collecting video) and could be helpful for consumers to see up front.

Our quantitative analysis indicated that device and cloud retention are among the four most influential attributes to decrease risk perception and increase willingness to purchase the device (see Table II). Participants strongly preferred no retention time over indefinite retention. The device and cloud retention attributes are currently placed on the secondary layer of the label [13]. Our quantitative findings suggest that these attributes should be promoted to the primary layer to better inform consumers' risk perception.

3) *Reducing Misconceptions*: As mentioned in Section IV-E, open-ended responses showed that some participants had misconceptions about the implications of patches and security updates, believing that the need for updates or patches indicates poor security. The root of the misconception was lack of knowledge about the necessity of receiving security updates and patching device vulnerabilities. We believe that an IoT label could be an effective way not only to inform consumers' purchase behavior, but also to educate them about privacy and security practices. On the secondary layer of the label in the expanded view, manufacturers can provide consumers with explanations as to why a device needs to receive security updates and why it needs to be patched. Besides, manufacturers

can discuss the potential consequences and risks of not having these practices.

In our study, we tested the efficacy of the attributes *security update* and *average time to patch*. Our quantitative and qualitative findings suggested that *average time to patch* was less useful to participants, especially because they had little knowledge of what time frame would be reasonable. This attribute was not recommended by experts to be included on the proposed label [13] and our results suggest that it should not be added to the label.

B. Need for Usable Controls

The effect on usability and a desire to have control over their device were two of the most common reasons participants gave for being less willing to purchase a device with the most protective attributes, such as multi-factor authentication and automatic updates (see Section IV-E).

From the quantitative analysis, we found that having control over three types of data practices would significantly decrease the perceived risk and increase the willingness to purchase the device (see Figure 2h). Although the majority of our participants indicated that automatic security updates would decrease risk, this information did not impact their willingness to purchase, mostly due to the lack of user control implied by the factor (based on qualitative analysis). Aligned with prior work [89], our participants preferred to know about the details of each update before allowing installation. Although having control was statistically favorable for some attributes, continuously asking for users' consent, on the other hand, could lead to usability challenges. For instance, participants indicated that asking the user to consent to data collection would decrease risk, but in their open-ended responses mentioned that it would also be a barrier to using the device (see Section IV-E).

Considering both usability and the desire for control, IoT manufacturers should provide users with choices about the control level they would like to have over their devices and provide convenient interfaces for exercising that control. Moreover, since the ability to control has been shown to decrease perceived risk [90], [91], IoT manufacturers need to clearly convey the potential risks and benefits of each of the offered choices to bridge the gap between the perceived risks and actual risks [16], [17].

C. Risk Acceptability and Willingness to Purchase

Similar to food nutrition labels, knowing what is healthier does not automatically translate into healthier behaviors [92]. Our findings from the two models of risk perception and willingness to purchase indicate that privacy and security attributes have a greater impact on risk perception than on willingness to purchase. From the open-ended responses, we found that participants consider multiple factors when making purchase decisions, including price, whether it provides functionality they need, convenience, and the desire to try new technology. In addition, based on regression results (see Table II), we found that participants who had the smart device themselves were significantly more willing to purchase a new

one of the same type (smart speaker or smart light bulb). All these factors could lead to risk acceptability and, therefore, lower the impact of privacy and security attributes in changing willingness to purchase.

Our quantitative analysis showed that the tested privacy and security attributes significantly change participants' risk perception. However, an effective label should be able to successfully impact consumers' desire to purchase the device as well. While our study provides some insights into the impact on willingness to purchase, future work is needed in which participants are provided with complete labels in a realistic setting so that they can consider label information alongside other purchase factors, including price, brand, and ratings.

D. On the Usefulness of Labels

Labels have been widely used to increase consumers' awareness in various domains, including energy efficiency [42] and nutrition [40]. However, the actual impact of labels on consumers' purchase behavior depends on various factors, some of which are associated with the consumers themselves, while others are related to the way information is presented on the label. These include personal factors such as level of knowledge and motivation in processing the label information [43], [93]–[95]. Heike et al. found that motivation to maintain a healthy diet and nutritional knowledge have a significant impact on the use of nutrition labels [43].

Food nutrition label research has shown that having prior nutritional knowledge significantly impacts people's use of nutrition labels [96]. From nutrition label research, we also know that knowledge predicts motivation [97] and motivation predicts knowledge [45]. In the realm of IoT, we hypothesize that providing knowledge to some consumers through IoT labels could initiate a virtuous cycle of knowledge and motivation, which could impact purchase behavior. Media reports have increased awareness among consumers about IoT devices' privacy and security practices, which could incentivize them to seek out more information when shopping for devices. In addition, IoT labels would also provide information to privacy advocates and journalists, who can then publicize this information and help inform the public about devices that stand out for having both exceptionally bad and exceptionally good privacy and security attributes.

E. Path to Label Adoption

Widespread adoption of IoT privacy and security labels would allow consumers to compare products based on their privacy and security attributes. In addition, once a critical mass of products have labels, consumers will look for them and may put more trust in those products that have them.

It is unclear whether sufficient incentives exist for voluntary label adoption. If large well-known manufacturers start putting labels on their products, smaller manufacturers may do so as well to try to increase their trust with consumers. Retailers could incentivize adoption by requiring manufacturers of products they sell to label their products, or even by promoting labeled products, for example by placing them at the top of

search results. However, it should be noted that past efforts to encourage industry disclosure schemes [98] and standardized privacy disclosures have faltered in the absence of regulatory mandates [99], as manufacturers lack incentives to adopt voluntary disclosures. In addition, manufacturers may fear that some disclosures could reduce interest in their products from potential customers.

Use of labels may be mandated by regulations or strongly encouraged through "safe harbor" provisions. In the US, law makers have proposed bills that include labels for IoT devices [10], [11]. Outside of the US, governments have started developing labeling schemes for IoT devices as well. Governments of the UK [7], Finland [100], and Singapore [101] are the forerunners in labeling their smart devices.

In addition to standardized labeling schemes, enforcement mechanisms are needed to ensure that there are consequences for companies that convey inaccurate information on their labels. In the US, the Federal Trade Commission or state attorneys general would likely prosecute companies who make false claims on their labels, similar to what happens when companies are found to make false claims in their privacy policies [8], [102], [103]. In addition, privacy and security rating or certification programs could provide seals to indicate they have independently verified the accuracy of label information [62], [104]–[106].

In the US, Underwriters Laboratories (UL) is assessing the security practices of IoT devices [77] and assigning a rating at one of five levels [62]. As more devices undergo assessment, these certifications may be added to IoT labels to further inform consumers' purchase decision making.

VI. CONCLUSION

Consumers are not aware of the privacy and security practices of their smart devices, and this lack of knowledge could expose them to privacy and security risks. One possible solution to better inform consumers' IoT device purchase decisions is to provide privacy and security information on a label, similar to a nutrition label for food. Little research has been done to test the efficacy of label information with IoT consumers.

We conducted an online study with 1,371 MTurk participants to measure information efficacy of 33 IoT label attribute-value pairs along two dimensions: risk perception and willingness to purchase. Overall, we found that label attributes successfully conveyed risk to participants, but we also identified some misconceptions. We found that label information more strongly influenced participants' risk perception than their willingness to purchase. Based on our findings, we propose recommendations to more effectively convey risks to IoT consumers.

ACKNOWLEDGMENTS

We thank our reviewers for their feedback, which helped improve this paper. This work was supported in part by DARPA and the Air Force Research Laboratory FA8750-15-2-0277, NSF awards TWC-1564009 and SaTC-1801472, and the Carnegie Mellon CyLab Security and Privacy Institute.

REFERENCES

- [1] Centre for International Governance Innovation & IPSOS, "2016 CIGI-IPSOS global survey on internet security and trust," <https://www.cigionline.org/internet-survey-2016>, 2016, accessed: 2020-06-03.
- [2] C. Thorun, M. Vetter, L. Reisch, and A. K. Zimmer, "Indicators of consumer protection and empowerment in the digital world," https://www.vzbv.de/sites/default/files/downloads/2017/03/13/conpolicy_executive_summary.pdf, March 2017, accessed: 2020-06-03.
- [3] V. Lara, "What the Internet of Things means for consumer privacy," <https://perspectives.eiu.com/technology-innovation/what-internet-things-means-consumer-privacy-0/white-paper/what-internet-things-means-consumer-privacy>, March 2018, accessed: 2020-06-03.
- [4] J. Phelps, G. Nowak, and E. Ferrell, "Privacy concerns and consumer willingness to provide personal information," *Journal of Public Policy & Marketing*, vol. 19, no. 1, pp. 27–41, 2000, accessed: 2020-06-03.
- [5] 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 CHI Conference on Human Factors in Computing Systems*. ACM, 2019, p. 534, accessed: 2020-06-03.
- [6] J. M. Blythe, N. Sombatrung, and S. D. Johnson, "What security features and crime prevention advice is communicated in consumer IoT device manuals and support pages?" *Journal of Cybersecurity*, vol. 5, no. 1, 06 2019, accessed: 2020-06-03. [Online]. Available: <https://doi.org/10.1093/cybsec/tyz005>
- [7] Department for Digital, Culture, Media and Sport, "Code of practice for consumer IoT security," <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security/code-of-practice-for-consumer-iot-security>, accessed: 2020-06-03.
- [8] Federal Trade Commission, "Comment to National Telecommunications and Information Administration," <https://www.ftc.gov/policy/advocacy/advocacy-filings/2017/06/ftc-comment-national-telecommunications-information>, June 2017, accessed: 2020-06-03.
- [9] The National Telecommunications and Information Administration, "Communicating IoT device security update capability to improve transparency for consumers," https://www.ntia.doc.gov/files/ntia/publications/communicating_iot_security_update_capability_for_consumers_-_jul_2017.pdf, accessed: 2020-06-03.
- [10] T. Lieu, "H.R.4163: Cyber Shield Act of 2017," <https://www.congress.gov/115/bills/hr4163/BILLS-115hr4163ih.pdf>, October 2017, accessed: 2020-06-03.
- [11] E. Markey, "S.2020: Cyber shield act of 2017," <https://www.congress.gov/115/bills/s2020/BILLS-115s2020is.pdf>, October 2017, accessed: 2020-06-03.
- [12] European Commission, "Proposal for a regulation of the European Parliament and of the Council on ENISA, the EU Cybersecurity Agency, and repealing regulation (EU) 526/2013, and on information and communication technology cybersecurity certification ("Cybersecurity Act")," <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52017PC0477&rid=1>, 2017, accessed: 2020-06-03.
- [13] P. Emami-Naeini, Y. Agarwal, L. Cranor, and H. Hibshi, "Ask the experts: What should be on an IoT privacy and security label?" in *IEEE Security & Privacy*. Los Alamitos, CA, USA: IEEE Computer Society, may 2020, pp. 771–788, accessed: 2020-06-03. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SP40000.2020.00043>
- [14] B. Wynne, "Institutional mythologies and dual societies in the management of risk," in *The Risk Analysis Controversy*. Springer, 1982, pp. 127–143, accessed: 2020-06-03.
- [15] L. Sjöberg, B.-E. Moen, and T. Rundmo, "Explaining risk perception. An evaluation of the psychometric paradigm in risk perception research," *Rotunde Publikasjoner Rotunde*, vol. 84, pp. 55–76, 2004, accessed: 2020-06-03.
- [16] B. Schneier, *Beyond fear: Thinking sensibly about security in an uncertain world*. Springer Science & Business Media, 2006, accessed: 2020-06-03.
- [17] M. M. Turner, C. Skubisz, and R. N. Rimal, "Theory and practice in risk communication: A review of the literature and visions for the future," in *The Routledge Handbook of Health Communication*. Routledge, 2011, pp. 174–192, accessed: 2020-06-03.
- [18] M. Harbach, S. Fahl, and M. Smith, "Who's afraid of which bad wolf? A survey of it security risk awareness," in *2014 IEEE 27th Computer Security Foundations Symposium*. IEEE, 2014, pp. 97–110, accessed: 2020-06-03.
- [19] E. Zeng, S. Mare, and F. Roesner, "End user security & privacy concerns with smart homes," in *Symposium on Usable Privacy and Security (SOUPS)*, 2017, accessed: 2020-06-03.
- [20] X. Bellekens, A. Hamilton, P. Seeam, K. Nieradzinska, Q. Franssen, and A. Seeam, "Pervasive ehealth services a security and privacy risk awareness survey," in *2016 International Conference on Cyber Situational Awareness, Data Analytics And Assessment (CyberSA)*. IEEE, 2016, pp. 1–4, accessed: 2020-06-03.
- [21] J. King and A. Selcukoglu, "Where's the beep? A case study of user misunderstandings of RFID," in *2011 IEEE International Conference on RFID*. IEEE, 2011, pp. 192–199, accessed: 2020-06-03.
- [22] V. Garg, K. Benton, and L. J. Camp, "The privacy paradox: A Facebook case study," in *2014 TPRC Conference Paper*, 2014, accessed: 2020-06-03.
- [23] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *IEEE Security & Privacy*, vol. 3, no. 1, pp. 26–33, 2005, accessed: 2020-06-03.
- [24] M. W. Skirpan, T. Yeh, and C. Fiesler, "What's at stake: Characterizing risk perceptions of emerging technologies," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–12, accessed: 2020-06-03.
- [25] A. Wieneke, C. Lehrer, R. Zeder, and R. Jung, "Privacy-related decision-making in the context of wearable use," in *PACIS*, 2016, p. 67, accessed: 2020-06-03.
- [26] S. Zheng, N. Aphorpe, M. Chetty, and N. Feamster, "User perceptions of smart home IoT privacy," *Proc. ACM Hum.-Comput. Interact.*, vol. 2, no. CSCW, Nov. 2018, accessed: 2020-06-03. [Online]. Available: <https://doi.org/10.1145/3274469>
- [27] I. Phau, M. Sequeira, and S. Dix, "Consumers' willingness to knowingly purchase counterfeit products," *Direct Marketing: An International Journal*, 2009, accessed: 2020-06-03.
- [28] V. A. Zeithaml, L. L. Berry, and A. Parasuraman, "The behavioral consequences of service quality," *Journal of Marketing*, vol. 60, no. 2, pp. 31–46, 1996, accessed: 2020-06-03.
- [29] I. Ajzen and M. Fishbein, *Understanding attitudes and predicting social behavior*. Prentice-Hall Englewood Cliffs, N.J, 1980, accessed: 2020-06-03.
- [30] R. L. Oliver and W. O. Bearden, "Crossover effects in the theory of reasoned action: A moderating influence attempt," *Journal of Consumer Research*, vol. 12, no. 3, pp. 324–340, 1985, accessed: 2020-06-03.
- [31] H. Karjalainen, J. Karvonen, M. Kesti, T. Koivumäki, M. Manninen, J. Pakola, A. Ristola, and J. Salo, "Factors affecting consumer choice of mobile phones: Two studies from Finland," *Journal of Euromarketing*, vol. 14, no. 3, pp. 59–82, 2005, accessed: 2020-06-03.
- [32] Z. Mack and S. Sharples, "The importance of usability in product choice: A mobile phone case study," *Ergonomics*, vol. 52, no. 12, pp. 1514–1528, 2009, accessed: 2020-06-03.
- [33] N. Saif, N. Razaq, M. Amad, and S. Gul, "Factors affecting consumers' choice of mobile phone selection in pakistan," *European Journal of Business and Management*, vol. 4, no. 12, pp. 16–26, 2012, accessed: 2020-06-03.
- [34] J. C. Olson and J. Jacoby, "Cue utilization in the quality perception process," *ACR Special Volumes*, 1972, accessed: 2020-06-03.
- [35] L. F. Cranor, J. Reagle, and M. S. Ackerman, "Beyond concern: Understanding net users' attitudes about online privacy," *The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy*, pp. 47–70, 2000, accessed: 2020-06-03.
- [36] H. Interactive, "A survey of consumer privacy attitudes and behaviors," *Rochester, NY*, vol. 47, 2000, accessed: 2020-06-03.
- [37] J. Turow, L. Feldman, and K. Meltzer, "Open to exploitation: America's shoppers online and offline," *Departmental Papers (ASC)*, p. 35, 2005, accessed: 2020-06-03.
- [38] J. Y. Tsai, S. Egelman, L. Cranor, and A. Acquisti, "The effect of online privacy information on purchasing behavior: An experimental study," *Information Systems Research*, vol. 22, no. 2, pp. 254–268, 2011, accessed: 2020-06-03.
- [39] P. G. Kelley, L. F. Cranor, and N. Sadeh, "Privacy as part of the app decision-making process," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, 2013, pp. 3393–3402, accessed: 2020-06-03.
- [40] Food and Drug Administration, "Nutrition facts label better informs your food choices," <https://www.fda.gov/ForConsumers/ConsumerUpdates/ucm387114.htm>, August 2016, accessed: 2020-06-03.

- [41] Federal Trade Commission, "Shopping for home appliances? use the energyguide label," <https://www.consumer.ftc.gov/articles/0072-shopping-home-appliances-use-energyguide-label>, January 2015, accessed: 2020-06-03.
- [42] European Union, "Energy efficient products," <https://ec.europa.eu/energy/en/topics/energy-efficiency/energy-efficient-products>, 2017, accessed: 2020-06-03.
- [43] S. Hieke and C. R. Taylor, "A critical review of the literature on nutritional labeling," *Journal of Consumer Affairs*, vol. 46, no. 1, pp. 120–156, 2012, accessed: 2020-06-03.
- [44] A. C. Drichoutis, R. M. Nayga, Jr, and P. Lazaridis, "Can nutritional label use influence body weight outcomes?" *Kyklos*, vol. 62, no. 4, pp. 500–525, 2009, accessed: 2020-06-03.
- [45] L. M. S. Miller and D. L. Cassady, "Making healthy food choices using nutrition facts panels: the roles of knowledge, motivation, dietary modifications goals, and age," *Appetite*, vol. 59, no. 1, pp. 129–139, 2012, accessed: 2020-06-03.
- [46] R. M. Nayga Jr, "Nutrition knowledge, gender, and food label use," *Journal of Consumer Affairs*, vol. 34, no. 1, pp. 97–112, 2000, accessed: 2020-06-03.
- [47] Apple, "App privacy details on the App Store," <https://developer.apple.com/app-store/app-privacy-details/>, accessed: 2021-3-8.
- [48] L. Tanczer, J. Blythe, F. Yahya, I. Brass, M. Elsdén, J. Blackstock, and M. Carr, "Summary literature review of industry recommendations and international developments on IoT security," *PETRAS IoT Hub, Department for Digital, Culture, Media & Sport (DCMS)*, 2018, accessed: 2020-06-03.
- [49] Symantec, "Why we need a security and privacy "nutrition label" for IoT devices," <https://www.symantec.com/blogs/expert-perspectives/why-we-need-security-and-privacy-nutrition-label-iot-devices>, February 2019, accessed: 2020-06-03.
- [50] British Standards Institution, "BSI launches kitemark for Internet of Things devices," <https://www.bsigroup.com/en-GB/about-bsi/media-centre/press-releases/2018/may/bsi-launches-kitemark-for-internet-of-things-devices/>, March 2018, accessed: 2020-06-03.
- [51] L. Sjöberg, "Factors in risk perception," *Risk Analysis*, vol. 20, no. 1, pp. 1–12, 2000, accessed: 2020-06-03.
- [52] H. Schütz and P. M. Wiedemann, "Judgments of personal and environmental risks of consumer products—do they differ?" *Risk Analysis*, vol. 18, no. 1, pp. 119–129, 1998, accessed: 2020-06-03.
- [53] N. Neil, P. Slovic, and P. Hakkinen, "Mapping consumer perceptions of risk," *Washington, DC: Chem. Manufactures Assoc*, 1993, accessed: 2020-06-03.
- [54] F. Farahmand and E. H. Spafford, "Understanding insiders: An analysis of risk-taking behavior," *Information Systems Frontiers*, vol. 15, no. 1, pp. 5–15, 2013, accessed: 2020-06-03.
- [55] V. Garg and J. Camp, "End user perception of online risk under uncertainty," in *2012 45th Hawaii International Conference on System Sciences*. IEEE, 2012, pp. 3278–3287, accessed: 2020-06-03.
- [56] L. Sjöberg and A. Biel, "Mood and belief-value correlation," *Acta Psychologica*, vol. 53, no. 3, pp. 253–270, 1983, accessed: 2020-06-03.
- [57] N. Fortes, P. Rita, and M. Pagani, "The effects of privacy concerns, perceived risk and trust on online purchasing behaviour," *International Journal of Internet Marketing and Advertising*, vol. 11, no. 4, pp. 307–329, 2017, accessed: 2020-06-03.
- [58] K. Backor, S. Golde, and N. Nie, "Estimating survey fatigue in time use study," in *International Association for Time Use Research Conference*. Washington, DC: Citeseer, 2007, accessed: 2020-06-03.
- [59] E. Dreyfuss, "A bot panic hits Amazon's Mechanical Turk," <https://www.wired.com/story/amazon-mechanical-turk-bot-panic/>, accessed: 2020-06-03.
- [60] Qualtrics, "Response quality," <https://www.qualtrics.com/support/survey-platform/survey-module/survey-checker/response-quality/>, accessed: 2020-06-03.
- [61] C. Olson, "New report tackles tough questions on voice and AI," <https://about.ads.microsoft.com/en-us/blog/post/april-2019/new-report-tackles-tough-questions-on-voice-and-ai>, accessed: 2020-06-03.
- [62] Underwriters Laboratories, "Identity management & security," <https://ims.ul.com/IoT-security-rating>, accessed: 2020-06-03.
- [63] The Digital Standard, "The standard," <https://www.thedigitalstandard.org/the-standard>.
- [64] ioXt, "The ioXt security pledge," <https://static1.squarespace.com/static/5c6dbac1f8135a29c7fbb621/t/5ca695ffec6eb0769f5608d1/1554421249364/ioXt-SecurityPledge-booklet-final.pdf>, accessed: 2020-06-03.
- [65] European Telecommunications Standards Institute, "Cyber security for consumer Internet of Things," https://www.etsi.org/deliver/etsi_ts/103600_103699/103645/01.01.01_60/ts_103645v010101p.pdf, accessed: 2020-06-03.
- [66] IoT Security Foundation, "IoT security compliance framework," <https://www.iotsecurityfoundation.org/wp-content/uploads/2018/12/IoTSF-IoT-Security-Compliance-Framework-Release-2.0-December-2018.pdf>.
- [67] Cellular Telecommunications Industry Association, "CTIA cybersecurity certification test plan for IoT devices," https://theinternetofthings.report/Resources/Whitepapers/853a8ac3-06c6-4975-9287-cb2403fa7617_CTIA-IoT-Test-Plan-V1_0.pdf, accessed: 2020-06-03.
- [68] R. H. B. Christensen, "A tutorial on fitting cumulative link mixed models with clmm2 from the ordinal package," *Vienna, Austria*, 2019, accessed: 2020-06-03.
- [69] J. Saldaña, *The coding manual for qualitative researchers*. Sage, 2015, accessed: 2020-06-03.
- [70] J. L. Fleiss, B. Levin, and M. C. Paik, *Statistical methods for rates and proportions*. John Wiley & Sons, 2013, accessed: 2020-06-03.
- [71] D. Von Winterfeldt, R. S. John, and K. Borcherting, "Cognitive components of risk ratings," *Risk Analysis*, vol. 1, no. 4, pp. 277–287, 1981, accessed: 2020-06-03.
- [72] R. Kang, S. Brown, L. Dabbish, and S. Kiesler, "Privacy attitudes of Mechanical Turk workers and the US public," in *10th Symposium on Usable Privacy and Security (SOUPS 2014)*, 2014, pp. 37–49, accessed: 2020-06-03.
- [73] United States Census Bureau, "Educational attainment in the United States: 2018," <https://www.census.gov/data/tables/2018/demo/education-attainment/cps-detailed-tables.html>, accessed: 2020-06-03.
- [74] H. Akoglu, "User's guide to correlation coefficients," *Turkish journal of emergency medicine*, vol. 18, no. 3, pp. 91–93, 2018.
- [75] K. P. Burnham and D. R. Anderson, "Multimodel inference: understanding aic and bic in model selection," *Sociological methods & research*, vol. 33, no. 2, pp. 261–304, 2004.
- [76] S. Mangiafico and M. S. Mangiafico, "Package 'rcompanion'," *Cran Repos*, pp. 1–71, 2017.
- [77] Underwriters Laboratories, "Methodology for marketing claim verification: Security capabilities verified to level bronze/silver/gold/platinum/diamond, UL MCV 1376," <https://www.shopulstandards.com/ProductDetail.aspx?UniqueKey=35953>, accessed: 2020-06-03.
- [78] E. U. Weber, N. Siebenmorgen, and M. Weber, "Communicating asset risk: How name recognition and the format of historic volatility information affect risk perception and investment decisions," *Risk Analysis: An International Journal*, vol. 25, no. 3, pp. 597–609, 2005, accessed: 2020-06-03.
- [79] J. Ginges, I. Hansen, and A. Norenzayan, "Religion and support for suicide attacks," *Psychological Science*, vol. 20, no. 2, pp. 224–230, 2009, accessed: 2020-06-03.
- [80] A. Boholm, "Comparative studies of risk perception: a review of twenty years of research," *Journal of Risk Research*, vol. 1, no. 2, pp. 135–163, 1998, accessed: 2020-06-03.
- [81] B. Richardson, J. Sorensen, and E. J. Soderstrom, "Explaining the social and psychological impacts of a nuclear power plant accident 1," *Journal of Applied Social Psychology*, vol. 17, no. 1, pp. 16–36, 1987, accessed: 2020-06-03.
- [82] G. Cvetkovich and P. L. Winter, "Trust and social representations of the management of threatened and endangered species," *Environment and Behavior*, vol. 35, no. 2, pp. 286–307, 2003, accessed: 2020-06-03.
- [83] T. C. Earle and G. Cvetkovich, "Social trust and culture in risk management," in *Social Trust and the Management of Risk*. Routledge, 2013, pp. 23–35, accessed: 2020-06-03.
- [84] A. Meijnders, C. Midden, A. Olofsson, S. Öhman, J. Matthes, O. Bondarenko, J. Gutteling, and M. Rusanen, "The role of similarity cues in the development of trust in sources of information about gm food," *Risk Analysis: An International Journal*, vol. 29, no. 8, pp. 1116–1128, 2009, accessed: 2020-06-03.
- [85] K. Nakayachi and G. Cvetkovich, "Public trust in government concerning tobacco control in japan," *Risk Analysis: An International Journal*, vol. 30, no. 1, pp. 143–152, 2010, accessed: 2020-06-03.
- [86] M. Siegrist, G. Cvetkovich, and C. Roth, "Salient value similarity, social trust, and risk/benefit perception," *Risk Analysis*, vol. 20, no. 3, pp. 353–362, 2000, accessed: 2020-06-03.

- [87] N. Allum, "An empirical test of competing theories of hazard-related trust: The case of gm food," *Risk Analysis: An International Journal*, vol. 27, no. 4, pp. 935–946, 2007, accessed: 2020-06-03.
- [88] European Union Agency for Cybersecurity, "Baseline security recommendations for IoT," <https://www.enisa.europa.eu/publications/baseline-security-recommendations-for-iot>, accessed: 2020-06-03.
- [89] K. Rose, S. Eldridge, and L. Chapin, "The Internet of Things: An overview," *The Internet Society (ISOC)*, vol. 80, 2015, accessed: 2020-06-03.
- [90] F. P. McKenna, "It won't happen to me: Unrealistic optimism or illusion of control?" *British Journal of Psychology*, vol. 84, no. 1, pp. 39–50, 1993, accessed: 2020-06-03.
- [91] L. E. Davis, A. Melmed, and R. Krop, *Individual preparedness and response to chemical, radiological, nuclear, and biological terrorist attacks*. Rand Corporation, 2003, accessed: 2020-06-03.
- [92] I. Ikonen, F. Sotgiu, A. Aydinli, and P. W. Verlegh, "Consumer effects of front-of-package nutrition labeling: An interdisciplinary meta-analysis," *Journal of the Academy of Marketing Science*, pp. 1–24, 2019, accessed: 2020-06-03.
- [93] L. R. Szykman, P. N. Bloom, and A. S. Levy, "A proposed model of the use of package claims and nutrition labels," *Journal of Public Policy & Marketing*, vol. 16, no. 2, pp. 228–241, 1997, accessed: 2020-06-03.
- [94] P. Klopp and M. MacDonald, "Nutrition labels: An exploratory study of consumer reasons for nonuse," *Journal of Consumer Affairs*, vol. 15, no. 2, pp. 301–316, 1981, accessed: 2020-06-03.
- [95] C. Moorman, K. Diehl, D. Brinberg, and B. Kidwell, "Subjective knowledge, search locations, and consumer choice," *Journal of Consumer Research*, vol. 31, no. 3, pp. 673–680, 2004, accessed: 2020-06-03.
- [96] L. M. S. Miller and D. L. Cassady, "The effects of nutrition knowledge on food label use. A review of the literature," *Appetite*, vol. 92, pp. 207–216, 2015, accessed: 2020-06-03.
- [97] L. M. S. Miller, T. N. Gibson, and E. A. Applegate, "Predictors of nutrition information comprehension in adulthood," *Patient Education and Counseling*, vol. 80, no. 1, pp. 107–112, 2010, accessed: 2020-06-03.
- [98] D. Esther, "Release 1.0," <http://cdn.oreilly.com/radar/r1/02-97.pdf>, accessed: 2020-10-24.
- [99] L. F. Cranor, "Necessary but not sufficient: Standardized mechanisms for privacy notice and choice," *J. on Telecomm. & High Tech. L.*, vol. 10, p. 273, 2012, accessed: 2020-06-03.
- [100] Finnish Transport and Communication Agency, "Finland becomes the first European country to certify safe smart devices – new cybersecurity label helps consumers buy safer products," <https://www.traficom.fi/en/news/finland-becomes-first-european-country-certify-safe-smart-devices-new-cybersecurity-label>.
- [101] Cyber Security Agency, "Cybersecurity labelling scheme," <https://www.csa.gov.sg/programmes/cybersecurity-labelling>.
- [102] Federal Trade Commission, "Lenovo, inc." <https://www.ftc.gov/enforcement/cases-proceedings/152-3134/lenovo-inc>, September 2017, accessed: 2020-06-03.
- [103] —, "Acidi group llc," <https://www.ftc.gov/enforcement/cases-proceedings/162-3103/acdi-group-llc>, June 2017, accessed: 2020-06-03.
- [104] S. Nielsen, "Should connected devices carry an IoT security-star rating?" <https://internetofthingsagenda.techtarget.com/blog/IoT-Agenda/Should-connected-devices-carry-an-IoT-security-star-rating>, May 2019, accessed: 2020-06-03.
- [105] R. Chirgwin, "Australia's IoT security rating might work, if done right," https://www.theregister.co.uk/2017/10/17/iot_security_rating_it_can_work_if_done_right/, October 2017, accessed: 2020-06-03.
- [106] U.S. Senate Committee on Commerce, Science, & Transportation, "Strengthening the cybersecurity of the Internet of Things," <https://www.commerce.senate.gov/public/index.cfm/hearings?ID=A6113AB7-E89B-48C7-B555-E3CBB1466040>, April 2019, accessed: 2020-06-03.

Security & Privacy Overview

Smart Security Camera, NS200
Firmware version 2.5.1: updated on: 6/15/2019
The device was manufactured in: United States

Casa

	<table style="width: 100%;"> <tr> <td style="width: 50%;">Security updates</td> <td>Automatic (available until 1/1/2022)</td> </tr> <tr> <td>Access control</td> <td>Password, Factory default, User-changeable, Multiple user accounts are allowed</td> </tr> </table>	Security updates	Automatic (available until 1/1/2022)	Access control	Password, Factory default, User-changeable, Multiple user accounts are allowed								
Security updates	Automatic (available until 1/1/2022)												
Access control	Password, Factory default, User-changeable, Multiple user accounts are allowed												
	<table style="width: 100%;"> <tr> <td style="width: 50%;">Sensor data collection</td> <td style="width: 50%; text-align: center;"> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> Video </div> <div style="text-align: center;"> Audio </div> </div> </td> </tr> <tr> <td>Purpose</td> <td>Providing device functions, research</td> </tr> <tr> <td>Data stored on device</td> <td>Identified</td> </tr> <tr> <td>Data stored on cloud</td> <td>Identified, Option to delete</td> </tr> <tr> <td>Shared with</td> <td>Manufacturer</td> </tr> <tr> <td>Sold to</td> <td>Not sold</td> </tr> </table>	Sensor data collection	<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> Video </div> <div style="text-align: center;"> Audio </div> </div>	Purpose	Providing device functions, research	Data stored on device	Identified	Data stored on cloud	Identified, Option to delete	Shared with	Manufacturer	Sold to	Not sold
Sensor data collection	<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> Video </div> <div style="text-align: center;"> Audio </div> </div>												
Purpose	Providing device functions, research												
Data stored on device	Identified												
Data stored on cloud	Identified, Option to delete												
Shared with	Manufacturer												
Sold to	Not sold												
	<div style="display: flex; justify-content: space-between;"> <div>Other collected data</div> <div>Presence, Temperature, Carbon monoxide, Usage information, User-entered information</div> </div>												
	<div style="display: flex; justify-content: space-between;"> <div>Privacy policy</div> <div>www.NS200.example.com/privacypolicy</div> </div>												

Detailed Security & Privacy Label:

www.iotsecurityprivacy.org/labels

More Information

Fig. 3: Primary layer of the label.

Security & Privacy Details

Smart Security Camera, NS200
Firmware version 2.5.1, updated on: 6/15/2019
The device was manufactured in: United States

Casa

	<table style="width: 100%;"> <tr> <td style="width: 50%;">Security updates</td> <td>Automatic (available until 1/1/2022)</td> </tr> <tr> <td>Access control</td> <td>Password, Factory default, User-changeable, Multiple user accounts are allowed</td> </tr> <tr> <td>Security oversight</td> <td>Audits performed by internal security auditors</td> </tr> <tr> <td>Ports and protocols</td> <td>www.NS200.example.com/port</td> </tr> <tr> <td>Hardware safety</td> <td>www.NS200.example.com/hwsafety</td> </tr> <tr> <td>Software safety</td> <td>www.NS200.example.com/swsafety</td> </tr> <tr> <td>Personal safety</td> <td>www.NS200.example.com/usersafety</td> </tr> <tr> <td>Vulnerability disclosure and management</td> <td>www.NS200.example.com/vulnreport</td> </tr> <tr> <td>Software and hardware composition list</td> <td>www.NS200.example.com/BOM</td> </tr> <tr> <td>Encryption and key management</td> <td>www.NS200.example.com/key</td> </tr> </table>	Security updates	Automatic (available until 1/1/2022)	Access control	Password, Factory default, User-changeable, Multiple user accounts are allowed	Security oversight	Audits performed by internal security auditors	Ports and protocols	www.NS200.example.com/port	Hardware safety	www.NS200.example.com/hwsafety	Software safety	www.NS200.example.com/swsafety	Personal safety	www.NS200.example.com/usersafety	Vulnerability disclosure and management	www.NS200.example.com/vulnreport	Software and hardware composition list	www.NS200.example.com/BOM	Encryption and key management	www.NS200.example.com/key																																								
Security updates	Automatic (available until 1/1/2022)																																																												
Access control	Password, Factory default, User-changeable, Multiple user accounts are allowed																																																												
Security oversight	Audits performed by internal security auditors																																																												
Ports and protocols	www.NS200.example.com/port																																																												
Hardware safety	www.NS200.example.com/hwsafety																																																												
Software safety	www.NS200.example.com/swsafety																																																												
Personal safety	www.NS200.example.com/usersafety																																																												
Vulnerability disclosure and management	www.NS200.example.com/vulnreport																																																												
Software and hardware composition list	www.NS200.example.com/BOM																																																												
Encryption and key management	www.NS200.example.com/key																																																												
	<table style="width: 100%;"> <tr> <td style="width: 15%;">Sensor data collection</td> <td style="width: 15%; text-align: center;"> Video </td> <td style="width: 15%; text-align: center;"> Audio </td> <td style="width: 15%; text-align: center;"> Presence </td> <td style="width: 15%; text-align: center;"> Temperature </td> <td style="width: 15%; text-align: center;"> Carbon Monoxide </td> </tr> <tr> <td>Collection frequency</td> <td>When user requests it</td> <td>Continuous</td> <td>Periodic, Option to opt-out</td> <td>Continuous, Option to opt-out</td> <td>Continuous, Option to opt-out</td> </tr> <tr> <td>Purpose</td> <td>Providing device functions, research</td> <td>Providing device functions, research</td> <td>Providing device functions</td> <td>Providing device functions</td> <td>Providing device functions</td> </tr> <tr> <td>Data stored on device</td> <td>Identified</td> <td>Identified</td> <td>Deidentified</td> <td>Deidentified</td> <td>Deidentified</td> </tr> <tr> <td>Local data retention time</td> <td>Up to a month</td> <td>Up to a month</td> <td>Up to a year</td> <td>Up to a year</td> <td>Up to a year</td> </tr> <tr> <td>Data stored on cloud</td> <td>Identified, Option to delete</td> <td>Identified, Option to delete</td> <td>No cloud storage</td> <td>Deidentified</td> <td>Deidentified</td> </tr> <tr> <td>Cloud data retention time</td> <td>Up to a month</td> <td>Up to a month</td> <td>Up to a month</td> <td>Up to a month</td> <td>Up to a month</td> </tr> <tr> <td>Shared with</td> <td>Manufacturer</td> <td>Manufacturer</td> <td>Not shared</td> <td>Manufacturer, Thirdparty</td> <td>Thirdparty, option to opt-out</td> </tr> <tr> <td>Sharing frequency</td> <td>Periodic, Adjustable</td> <td>Periodic, Adjustable</td> <td>Not shared</td> <td>Continuous</td> <td>Continuous</td> </tr> <tr> <td>Sold to</td> <td>Not sold</td> <td>Not sold</td> <td>Not sold</td> <td>Thirdparty</td> <td>Thirdparty, Option to opt-out</td> </tr> </table>	Sensor data collection	 Video	 Audio	 Presence	 Temperature	 Carbon Monoxide	Collection frequency	When user requests it	Continuous	Periodic, Option to opt-out	Continuous, Option to opt-out	Continuous, Option to opt-out	Purpose	Providing device functions, research	Providing device functions, research	Providing device functions	Providing device functions	Providing device functions	Data stored on device	Identified	Identified	Deidentified	Deidentified	Deidentified	Local data retention time	Up to a month	Up to a month	Up to a year	Up to a year	Up to a year	Data stored on cloud	Identified, Option to delete	Identified, Option to delete	No cloud storage	Deidentified	Deidentified	Cloud data retention time	Up to a month	Up to a month	Up to a month	Up to a month	Up to a month	Shared with	Manufacturer	Manufacturer	Not shared	Manufacturer, Thirdparty	Thirdparty, option to opt-out	Sharing frequency	Periodic, Adjustable	Periodic, Adjustable	Not shared	Continuous	Continuous	Sold to	Not sold	Not sold	Not sold	Thirdparty	Thirdparty, Option to opt-out
Sensor data collection	 Video	 Audio	 Presence	 Temperature	 Carbon Monoxide																																																								
Collection frequency	When user requests it	Continuous	Periodic, Option to opt-out	Continuous, Option to opt-out	Continuous, Option to opt-out																																																								
Purpose	Providing device functions, research	Providing device functions, research	Providing device functions	Providing device functions	Providing device functions																																																								
Data stored on device	Identified	Identified	Deidentified	Deidentified	Deidentified																																																								
Local data retention time	Up to a month	Up to a month	Up to a year	Up to a year	Up to a year																																																								
Data stored on cloud	Identified, Option to delete	Identified, Option to delete	No cloud storage	Deidentified	Deidentified																																																								
Cloud data retention time	Up to a month	Up to a month	Up to a month	Up to a month	Up to a month																																																								
Shared with	Manufacturer	Manufacturer	Not shared	Manufacturer, Thirdparty	Thirdparty, option to opt-out																																																								
Sharing frequency	Periodic, Adjustable	Periodic, Adjustable	Not shared	Continuous	Continuous																																																								
Sold to	Not sold	Not sold	Not sold	Thirdparty	Thirdparty, Option to opt-out																																																								
	<div style="display: flex; justify-content: space-between;"> <div>Other collected data</div> <div>Usage information, User-entered information</div> </div>																																																												
	<div style="display: flex; justify-content: space-between;"> <div>Data linkage</div> <div>Data may be linked with internal and external data sources</div> </div>																																																												
	What could be inferred from user's data	No data inference																																																											
	Special data handling practices for children	Yes																																																											
	In compliance with	GDPR, ISO27001																																																											
	Privacy policy	www.NS200.example.com/privacypolicy																																																											

Call Casa with your questions at

412-313-2793 (24/7 support)

More Information

Fig. 4: Secondary layer of the label.

APPENDIX B: SURVEY QUESTIONS

The following questions are for one of the experimental conditions, where the between-subject factors of *device_type* and *device_recipient* are smart light bulb and purchasing the device for a friend, respectively.

A. Device-Related Questions

Q1: How concerned are you about the way *smart light bulbs* with *presence sensors* collect, store, and use information related to *whether someone is present in the room*? (Answered on a 5-point scale from "Not at all concerned" to "Very concerned")
[If the answer to Q1 is "Only slightly concerned," "Somewhat concerned," "Moderately concerned," or "Very concerned"] **Q2:** What about data collection, storage, and use by smart light bulbs with *presence sensors* makes you feel concerned? [text entry]
[If the answer to Q1 is "Not at all concerned"] **Q2:** What about data collection, storage, and use by *smart light bulbs* with *presence sensors* makes you feel not at all concerned? [text entry]
Q3: Do you currently have a *smart light bulb* with *presence sensor* in your home? (Choices: "Yes," "No")
[If the answer to Q3 is "Yes"] **Q4:** How long have you had your *smart light bulb* with *presence sensor*? If you have more than one device, answer the question for the one that you have had for the longest time. (Choices: "Less than a month," "Between a month and a year," "More than a year," "I don't remember")
[If the answer to Q3 is "Yes"] **Q5:** How did you acquire your *smart light bulb* with *presence sensor*? (check as many as apply) (Choices: "I purchased it," "Somebody else in my home purchased it," "I received it as a gift," "It was installed by my landlord," "other (please specify)")
[If the answer to Q5 is "I purchased it"] **Q6:** What brand(s) of *smart light bulb* with *presence sensor* did you purchase? [text entry]
[If the answer to Q5 is "I purchased it"] **Q7:** What were your reasons to purchase a *smart light bulb* with *presence sensor*? [text entry]
[If the answer to Q3 is "No"] **Q4:** Have you ever been in the market to purchase a *smart light bulb* with *presence sensor*? (Choices: "No," "Yes")
[If the answer to Q4 is "No"] **Q5:** What made you decide not to purchase the *smart light bulb* with *presence sensor* that you were in the market for? [text entry]

B. Label-Related Questions

If one of the randomly-assigned *attribute-value* pairs was *security update:automatic*:

Imagine you are making a decision to purchase a *smart light bulb* for a *friend*. This device has a *presence sensor* that will *detect whether someone is present in the room to control the lighting automatically*. The price of the device is within your budget and the features are all what you would expect from a smart light bulb with *presence sensor*. On the package of the device there is a label that explains the privacy and security practices of the *smart light bulb* with *presence sensor*. The label on the device indicates the following:

Security update: Automatic

Q1: How confident are you that you know what this information on the label means? (Answered on a 5-point scale from "Not at all confident" to "Very confident")

1) Risk Perception:

Q2: I believe receiving *automatic security updates* (Choices: "Strongly decreases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*," "Slightly decreases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*," "Does not have any impact on the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*," "Slightly increases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*," or "Strongly increases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*")
[If the answer to Q2 is "Strongly decreases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*" or "Slightly decreases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*"] **Q3:** Please explain why you believe receiving *automatic security updates* decreases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*. [text entry]
[If the answer to Q2 is "Strongly increases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*" or "Slightly increases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*"] **Q3:** Please explain why you believe receiving *automatic security updates* increases the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*. [text entry]
[If the answer to Q2 is "Does not have any impact on the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*"] **Q3:** Please explain why you believe receiving *automatic security updates* does not have any impact on the privacy and security risks associated with this specific *smart light bulb* with *presence sensor*. [text entry]

2) Willingness to Purchase:

Q4: Assuming you were in the market to purchase this *smart light bulb* with *presence sensor* for a *friend as a gift*, knowing that the device will *automatically receive security updates*, would (Choices: "Strongly decrease your willingness to purchase this device for a *friend as a gift*," "Slightly decrease your willingness to purchase this device for a *friend as a gift*," "Not have any impact on your willingness to purchase this device for a *friend as a gift*," "Slightly increase your willingness to purchase this device for a *friend as a gift*," or "Strongly increase your willingness to purchase this device for a *friend as a gift*")
[If the answer to Q4 is "Strongly decrease your willingness to purchase this device for a *friend as a gift*" or "Slightly decrease your willingness to purchase this device for a *friend as a gift*"]
Q5: Please explain why knowing that the device will *automatically receive security updates* would decrease your willingness to purchase the device for a *friend as a gift*. [text entry]
[If the answer to Q4 is "Strongly increase your willingness to purchase this device for a *friend as a gift*" or "Slightly increase your willingness to purchase this device for a *friend as a gift*"]
Q5: Please explain why knowing that the device will *automatically receive security updates* would increase your willingness to purchase the device for a *friend as a gift*. [text entry]
[If the answer to Q4 is "Not have any impact on your willingness to purchase this device for a *friend as a gift*"] **Q5:** Please explain why knowing that the device will *automatically receive security updates* would not have any impact on your willingness to purchase the device for a *friend as a gift*. [text entry]

3) Attention-Check Question for Automatic Update:

Q6: Which statement is correct about the device described in the previous question? (Choices: "The device will automatically get updated," "The device will manually get updated," "The device will not get updated," or "The device will ask for my permission each time to install security updates.")

C. Functionality Perception

[When the device type is smart light bulb] **Q1:** How do you think a *smart light bulb* with *presence sensor* works? (Choices: "The device always senses whether someone is present in the room," "The device starts sensing whether someone is present in the room when you press a button to turn on the presence sensor on the device," "The device starts sensing whether someone is present in the room when you turn on the lights," or "I have no idea how a smart light bulb with presence sensor works.")
[When the device type is smart speaker] **Q1:** How do you think a *smart speaker* with *voice assistant* works? (Choices: "The device always listens to your voice to respond to your commands," "The device waits for you to mention the wake word (e.g., 'Alexa,' 'OK Google')," "The device starts listening when you press a button to turn on the microphone on the device," or "I have no idea how a smart speaker with voice assistant works.")

D. Demographic Questions

Q1: What is your age? [text entry]
Q2: What is your gender? [text entry]
Q3: What is the highest degree you have earned? (Choices: "No high school degree," "High school degree," "College degree," "Professional degree (Master's/PhD/medical/law)," "Associate degree," or "Prefer not to answer")
Q4: Do you have a background in technology (if yes, please specify what your background is)? [text entry]

APPENDIX C: CONSUMER EXPLANATIONS FOR ATTRIBUTE-VALUE PAIRS

Attribute-value	Consumer explanation
Security update: automatic	Device will automatically receive security updates
Security update: none	Device will not receive any security updates
Access control: multi-factor authentication	At least two independent factors to authenticate a user are required to access the device, for example a password and a confirmation from a previously registered phone
Access control: none	Anyone can access the device without a password or other authentication method
Purpose: device function	Data is being collected to provide the main device features, improve services, and help develop new features
Purpose: monetization	The manufacturer and service provider receive income from showing personalized advertisements to users or selling user's data to third parties
Device storage: none	The collected data will not be stored on the device
Device storage: identified	User's identity could be revealed from the data stored on the device
Cloud storage: none	The collected data will not be stored on the cloud
Cloud storage: identified	User's identity could be revealed from the data stored on the cloud
Shared with: none	Data is not being shared
Shared with: third parties	Data is being shared with third parties
Sold to: none	Data is not being sold
Sold to: third parties	Data is being sold to third parties
Average time to patch: 1 month	Vulnerabilities will be patched within 1 month of discovery
Average time to patch: 6 months	Vulnerabilities will be patched within 6 months of discovery
Security audit: internal & external	Security audits are performed by internal and third-party security auditors
Security audit: none	No security audit is being conducted
Collection frequency: on user demand	Data is collected when the user requests it
Collection frequency: continuous	When the device is turned on, it will continuously collect data until it is turned off
Sharing frequency: on user demand	Data is shared when the user requests it
Sharing frequency: continuous	When the device is turned on, it will continuously share data until it is turned off
Device retention: none	User's data will not be retained on the device
Device retention: indefinite	User's data may be retained on the device indefinitely
Cloud retention: none	User's data will not be retained on the cloud
Cloud retention: indefinite	User's data may be retained on the cloud indefinitely
Data linkage: none	Data will not be linked with other data sources
Data linkage: internal & external	Data may be linked with other information collected by the manufacturer as well as other information
Inference: none	No additional information about user will be inferred from user's data
Inference: additional information	User's characteristics and psychological traits, attitudes and preferences, aptitudes and abilities, and behaviors could be inferred from the collected data
Control over: cloud data deletion	User has an option to delete the data that is being stored on the cloud
Control over: device storage	Data will not be stored on the device unless user opts in to device storage
Control over: device retention	User can change the duration for which their data may be retained on the device

TABLE III: Consumer explanations that we presented for attribute-value pairs in the survey.

APPENDIX D: SURVEY DEMOGRAPHIC INFORMATION

Metric	Levels	MTurk (%)	Census (%)
Gender	Male	54.0	48.5
	Female	45.5	51.5
	Non-binary	0.5	—
Age	18-29 years	23.4	21.0
	30-49 years	61.3	33.4
	50-64 years	13.0	25.2
	65+ years	2.3	20.4
Education	No high school	0.3	10.9
	High school	28.9	47.2
	College	51.2	20.6
	Professional	10.6	11.7
	Associate	8.7	9.6
	No answer	0.3	—
Tech Background	Yes	19.8	—
	No	80.2	—

TABLE IV: Demographic information of our participants and 2018 US Census data [73]. In some cases, the Census data did not include a specific category, denoted by —.