



“I would still use it but I wouldn’t trust it”: Evaluating Mechanisms for Transparency and Control for Smart-Home Sensors

BEN WEINSHEL, Carnegie Mellon University, USA

YUVRAJ AGARWAL, Carnegie Mellon University, USA

LUJO BAUER, Carnegie Mellon University, USA

Smart-home devices, such as smart speakers and cameras, provide convenient home automation and media control, but the sensors that continuously collect data in users’ homes create privacy concerns. To attempt to increase trust, device manufacturers and researchers have developed privacy features, such as indicator lights, hardware controls, and microphone jammers. To inform the design of more trustworthy products, we conducted a 489-participant online survey to understand how device type, brand, and privacy features impact trust. Our survey also examined whether providing more information about privacy features’ limitations changed participants’ perceptions. Contrary to our expectations, device brand did not significantly impact trust. Hardware mute controls were most effective at increasing trust. Participants expressed high intent to use familiar software-backed features, while expressing reservations about novel features proposed by researchers (e.g., jamming devices). Participants’ reactions after seeing information about privacy features’ limitations varied by feature, suggesting that the features’ strengths and weaknesses are not equally well-understood. Based on our findings, we make several recommendations, including that device manufacturers and researchers explore making software-backed features more secure, as our results suggest that users may use those features even if they do not consider them reliable or trustworthy.

CCS Concepts: • **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → **Ubiquitous and mobile devices**.

Additional Key Words and Phrases: smart home, internet of things, privacy, security, sensor control

ACM Reference Format:

Ben Weinschel, Yuvraj Agarwal, and Lujo Bauer. 2025. “I would still use it but I wouldn’t trust it”: Evaluating Mechanisms for Transparency and Control for Smart-Home Sensors. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 2, Article 54 (June 2025), 33 pages. <https://doi.org/10.1145/3729480>

1 Introduction

Always-on smart-home devices with cameras and microphones have the potential to support new computational use cases such as improved comfort, convenience, and reduced energy consumption. Smart speakers and cameras have been widely adopted and enable new interactions with technology, yet persistent privacy concerns remain, including unwanted recording and unintentional data sharing [21, 36, 70]. Many people are not comfortable purchasing or using these products [10, 30, 41].

To address these concerns, device manufacturers have built privacy features into their products, such as hardware mute switches that disable the microphones and cameras [8, 33]. Simultaneously, the research community has proposed privacy features that could be built into products, such as loudness and gaze thresholds before triggering a voice assistant [49], and developed add-on accessories to provide additional protection, such as

Authors’ Contact Information: Ben Weinschel, bweinschel@cmu.edu, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA; Yuvraj Agarwal, yuvraj@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA; Lujo Bauer, lbauer@cmu.edu, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.



This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2474-9567/2025/6-ART54

<https://doi.org/10.1145/3729480>

ultrasound jammers [19, 51, 55, 64] and tangible controls [22, 23]. In limited experimental evaluations, these systems increased user comfort and trust [22, 49].

However, it remains unclear how the various proposed and existing privacy features compare with each other in terms of users' likelihood to use them in practice, the perceived reliability of the features, and how they impact trust in the smart-home product. Furthermore, users' reactions may vary based on limited understanding of the limitations and tradeoffs of the features, and prior work has not presented these details in a consistent manner. Some of the most privacy-protective solutions, including hardware mute switches, disable hands-free activation [8], which arguably breaks one of the core use cases for voice assistants. Prior work has found that product reviewers consider Apple's smart-home products to be more private and secure [34], while Amazon has had a number of privacy-related controversies [21]. However, prior work has not explored in detail whether intrinsic factors such as the device type or brand affect the perception of these privacy features.

We conducted a 489-participant online survey to more accurately assess the likely benefit of current and proposed privacy features for smart-home sensor products and inform future privacy features. Our study design is based around two research questions:

- **RQ1:** How do device type, brand, and information about privacy features contribute to users' intent to use the privacy features, the perceived reliability of the features, and trust in the smart-home sensor product?
- **RQ2:** How does information about the limitations of a privacy feature change users' perceptions of the feature and the smart-home product?

We found that participants expressed general distrust of *all* smart-home devices and their manufacturers. In many cases, the concerns users raised were not even smart-home specific, indicating a general attitude of privacy resignation. However, echoing prior work, participants who owned a similar type of smart-home product (e.g., a smart speaker) expressed higher trust in the products described in the survey [20]. In terms of the privacy features themselves, participants perceived hardware-backed controls to be the most reliable, and a hardware mute control increased trust in the smart-home product more than any other privacy features we evaluated. However, perceived reliability and trust were somewhat independent of participants' intent to use the features. Participants expressed similar intent to use mute controls (both hardware- and software-backed), indicator lights, and activity logs, despite the fact that they considered hardware mute controls significantly more reliable and that such controls engendered more trust in the smart-home product. Our results suggest that reliability of privacy features matters, but only to an extent—participants were relatively unconcerned about software bugs or vulnerabilities, but expressed more concern about features they considered difficult to test and verify.

After seeing information about the privacy features' limitations, participants indicated that they would be less likely to use most features, would consider them less reliable, and would trust the device less. This effect was strongest for the activity history screen and indicator light. These trends suggest that some reliability concerns, such as potential for software bugs, were not obvious to participants. However, more participants indicated they would use software-backed features than found them reliable, suggesting that users may use a software-backed privacy feature despite knowing it may not work as expected.

Our results highlight the need for privacy features to be easy to test and verify by both consumers and third-parties. Participants expressed they would use features that worked for the use cases they cared about and they were able to confirm were working as expected with minimal testing. However, also participants indicated they would use some features that are not guaranteed to be fully reliable (e.g., the software mute control), suggesting that it may be feasible for unscrupulous device vendors to trick people into relying on privacy features that should not be trusted.

2 Related Work

We discuss work on trust in IoT devices, smart-home sensor products, and sensor transparency and control.

2.1 Trust in IoT Devices

Many people express distrust in smart-home devices and their manufacturers, often based on prior experiences and perceptions about how their data is used outside of the IoT [66]. Privacy perceptions are often contextual depending on the scenario [27] and practices of the manufacturer [28]. Users who own a smart-home product often express higher trust in other smart-home products [20, 73], but general privacy attitudes do not differ between owners of different smart speaker brands [2]. Security, product performance, and brand can lead to trust in consumer products [15, 31, 32, 35, 43, 47, 50]. In our study, we build on this body of work by evaluating trust with respect to a variety of privacy features in smart-home sensor products.

2.2 Smart-Home Sensor Products

In this work, we focus on *smart-home sensor products*, where high-fidelity sensing is core to the product's functionality, such as smart speakers, displays, and cameras. Smart speakers, including the Amazon Echo and Apple HomePod, have been broadly adopted, yet they have caused substantial discussion about their privacy impacts. Discussion in the media has brought attention to these devices' always-listening behavior, potential for false triggers, and data-annotation practices [21, 36, 70]. In response to such concerns, companies have added additional transparency and control [11, 33]. Researchers have found that people express a variety of concerns, including about how the device is "always listening" [10], data-retention policies [45], and sharing of data with third-party apps or "skills" [1]. Existing privacy controls do not meet user needs [41].

Future voice interfaces may be always-listening [65], proactive [16], or invisible [42]. While in this paper we focus on existing products, other work has explored privacy features for always-listening assistants [46]. Other smart-home products including including cameras [73] and toys [48] present similar privacy risks.

Radio-frequency sensing technologies, such as millimeter wave, can reduce privacy exposure, as the sensors may reveal less data than a camera or microphone [56]. However, RF sensors can make inferences about activity, health conditions, and other sensitive data [52]. In many cases, the inferences made by these sensors impact privacy perceptions more than the (often-unintelligible) raw data [63]. Lower-fidelity sensors such as temperature and motion sensors may also reveal activity in the home [20], though prior work has found that users consider those sensors less sensitive [27]. In this work, we focus on cameras and microphones on existing smart-home products as people are likely have an intuitive understanding of what data cameras and microphones can capture, consider audio-visual data to be sensitive, and be familiar with smart speakers and similar products.

2.3 Sensor Transparency and Control

In this section, we discuss some of the privacy features that have been developed by both device manufacturers and researchers. In our study, we evaluate a selection of these features.

2.3.1 Transparency Mechanisms.

Indicator light: Commercially-available smart speakers generally include an indicator light that shows when the device is actively listening. Some devices, such as the Amazon Echo, have a light to indicate that the microphones are disabled [8]. Researchers evaluating laptop indicator lights have found many users do not notice them [58], and as such have explored making the indication more detailed through dynamically-displayed icons [24].

Activity log: Many devices offer an activity log, where the user can review a history of interactions, listen to audio recordings, and read transcripts of requests [8, 33]. This is advertised as a privacy feature that can build trust that the device is not recording unexpectedly [10]. However, the presence of an activity log introduces privacy tradeoffs: current implementations require server-side data storage, and the logs allow users to view others' activity, which has been used for surveillance [41].

Other interfaces: "Test drives," proposed by Malkin et al., allow users to audit apps running on voice assistants that are actively listening to all conversations (as opposed to listening for a trigger word) [44, 46].

Hidden device discovery: Researchers have additionally explored novel mechanisms for discovering hidden cameras and microphones through mechanisms such as Wi-Fi and RF signatures [62, 74].

2.3.2 Physical Controls.

Mute button: User studies have found that a microphone mute button that disables the microphones provides a sense of control [10], but some do not trust it to work as described [41]. In one study, only 5% of smart speaker owners reported using the mute button [45].

Tangible controls: Ahmad et al. introduced the notion of “tangible privacy” to provide stronger perceived privacy for bystanders, and contributed design recommendations based on an interview study [6]. In a separate work, Ahmad et al. evaluated the effectiveness of tangible controls by showing survey participants a smart display with various types of controls and feedback mechanisms, finding that physical controls increase perceptions of trust, reliability, usability, and control, but the feedback mechanism did not have a significant impact [5].

2.3.3 Add-On Protections.

Wireless microphone accessory: Do et al. developed a wireless microphone accessory, powered by RF backscatter, that provides a perceptible assurance power is disconnected when the clamshell-style accessory is closed [22].

Smart webcam cover: Additionally, in a separate work, Do et al. developed a smart webcam cover that automatically covers the camera when it is inactive, but can only uncover the camera with user interaction [23].

Disconnecting power: Chandrasekaran et al. evaluated a remote-controlled smart plug as a mechanism to provide users with an alternative control to a smart speaker’s own mute button [17]. Users have also been observed to adopt this privacy-protective behavior. Jin et al. surveyed users’ smart-home privacy-protective behaviors; one participant used smart plugs to turn their cameras on and off [38].

2.3.4 Jamming Devices. The research community has proposed other types of privacy features that block or jam microphones. Microphones can be jammed through ultrasound that is inaudible to humans, but due to the non-linearities of microphone hardware, creates signals in the audible range that block out actual sounds [60].

Jammer wristband: Ultrasound jamming can have varying effectiveness based on the exact positioning of the jammer and the microphone being blocked. To address this limitation, Chen et al. prototyped a wristband that, by broadcasting in many directions, can block multiple microphones more reliably [19, 37].

Jammer accessory: Several projects propose placing the jamming device on top of or next to the voice assistant. Chandrasekaran et al. used a remote accessory to control the jammer [17], while others included a microphone in the accessory, unjamming the smart speaker when a wake word [40, 55, 64] or clap [51, 54] is detected.

2.3.5 System Hardening. The risk of unintentional recording, or false triggers could be reduced by adding additional requirements or checks before triggering a voice assistant.

Interpersonal communication cues: Mhaidli et al. built a voice assistant that would only be activated when the user is looking at the device or speaking with a louder volume than normal conversation, and found that these controls were easier to use than existing hardware mute controls [49].

System security: De Vaere et al. used trusted execution environments to provide a hardened pipeline for wake word detection that would allow auditing when the device was activated [69].

Audio filtering: Other work has explored ways that audio can be processed to anonymize recordings [7, 61, 72].

In contrast to prior work that generally evaluates one or two privacy features, we evaluate a larger set of privacy features shipping in current projects and proposed by researchers. Additionally, while prior work has assessed user reactions to broad set of privacy and security factors [28], we are not aware of any work that has specifically looked at whether brand and device type impact perceptions of privacy features; hence, we evaluated those factors in our study. We also extend prior work by exploring how people evaluate the tradeoffs between stronger security guarantees and usability with respect to privacy feature with respect to privacy features.

Table 1. We measured participants' reactions using both Likert scales (this table) and free-response questions (Appendix A). Device, brand, and privacy feature varied.

(a) Baseline privacy concerns and trust in the device presented.		(b) Reactions to privacy features. The text in [] varied based on whether we had described the feature's limitations.	
Concern-DeviceType	Which of these choices best describes how you feel about how smart speaker products in general collect, store and use information?	WouldUse-Feature	[Knowing these limitations,] I would use the indicator light with a Apple smart speaker.
Concern-Brand	Which of these choices best describes how you feel about how Apple products collect, store and use information?	Reliable-Feature	[Knowing these limitations,] I trust that the indicator light would work reliably .
Trust-Collect	I would trust this Apple smart speaker to only collect data (e.g. audio recordings) when appropriate .	Trust-Device	Given [I could use the indicator light with this product the limitations of using the indicator light], I would trust that this Apple smart speaker protects my privacy .
Trust-Use	I would trust this Apple smart speaker to use and share data it collects appropriately .		
Trust-Device	I would trust this Apple smart speaker to protect my privacy .		

3 Methodology

We recruited participants for our survey through Prolific. Participants were required to be in the United States. We limited recruitment to participants who owned at least one smart-home product (e.g., smart TV, smart speaker, or smart lighting) so that they would be familiar with the concepts in our survey. Our survey took a median of 15 minutes. Participants were paid \$3 for completing the survey. The study was approved by our institution's IRB. The full survey instrument is included in Appendix A.

After reviewing the consent form and agreeing to participate in our survey, participants were shown information about a randomly-assigned device (described in Section 3.1.1) and asked about the device's privacy practices (Table 1a). Participants were then shown information about a randomly-selected privacy feature and answered questions about the feature (described in Section 3.1.2). To capture perceptions of the feature when provided with a basic description of its functionality, we asked participants about their willingness to use the feature, their perceived reliability of the feature, and their trust in the device's privacy practices, considering they can use the feature (Table 1b). We then provided participants with more detail about the potential limitations of the feature (described in Section 3.1.3), and asked the same questions again to measure the extent to which their reactions changed based on the new information. We also asked participants open-ended questions to learn about why their opinions did or did not change. The section about privacy features and their limitations was shown three times to each participant, with a different feature shown each time.

3.1 Experimental Conditions

Our study had three experimental variables: the device presented to participants, privacy features, and information about limitations. In this section, we discuss our choice of experimental conditions, the descriptions we presented to participants, and the between- or within-subjects design for each variable.

3.1.1 Devices. Participants were assigned to one of six device conditions: Amazon, Apple, Google, or Sustios smart speaker; Sustios smart display; or Sustios smart indoor camera. We were interested in understanding whether reactions to privacy features varied by the type of device, so we selected widely-available products that include microphones and cameras, and excluded products with infrared, temperature, and RF sensors for which users may not fully understand inferences that can be made from the collected data [63]. Our hypothesis was that brand reputation would impact perceptions of the privacy features (e.g., if participants trusted a brand, they might not see a need for additional privacy features) and so we varied the device brand, including three

Table 2. We selected privacy features that include widely-available features and those proposed by researchers.

Privacy feature	Category	Source
Software mute control	Software-backed features	Existing products [8, 12, 33]
Indicator light	Software-backed features	Existing products [8, 33]
Activity history screen	Software-backed features	Existing products [8, 33]
Hardware mute control	Hardware-backed features	Existing products [8, 33]
Wireless microphone	Hardware-backed features	Research prototypes [22]
Loudness detection feature	Audio heuristics/jamming	Research prototypes [49]
Jammer accessory	Audio heuristics/jamming	Research prototypes [17, 40, 51, 55, 64]
Jammer wristband	Audio heuristics/jamming	Research prototypes [19]

well-known brands (Amazon, Apple, and Google) and one fictional one (Sustios). We used a between-subjects design for device type and brand so participants could consider a single product for the duration of the survey.

At the beginning of the survey, participants were shown a description of the device that was selected for them and its features. We reviewed marketing materials for commercially-available products and selected features advertised by multiple companies. As we were looking to measure participants' perceptions of different device types and brand reputations, not specific features of a single brand's products, we wrote the descriptions to be similar across all three device types and identical between brands. For example, we included sound detection and smart-home integration as features on all products. Appendix B includes the full descriptions.

3.1.2 Privacy Features. We selected privacy features that encompass a variety of transparency and control mechanisms included with existing smart-home products and proposed by researchers based on a review of the recent literature and product marketing materials (Table 2). We reviewed the marketing materials for Amazon [8], Apple [12], and Google [33] smart speakers and displays and included the features that provide sensor transparency and control (as opposed to, e.g., controls for audio grading or analytics): hardware mute control, software mute control, indicator light, and activity history screen.

Our review of the academic literature started with work on a jammer wristband that received media attention [19, 37], and a 2023 publication on tangible microphone controls [22]. We then recursively reviewed the publications that cite and are cited by those papers, and from that review included a jammer accessory [17, 40, 51, 55, 64] and loudness detection feature [49] in our study (the concepts from multiple jammer accessories proposed were combined into a single feature). We additionally reviewed the proceedings of IMWUT, IEEE Symposium on Security & Privacy, USENIX Security Symposium, ACM CCS, and Privacy Enhancing Technologies Symposium from 2020 through 2024 to ensure we comprehensively included privacy features proposed by researchers. The features we studied have a variety of characteristics: some provide only transparency and others add control, and some have hardware-backed guarantees while others use software implementations.

For all the features, we reviewed existing literature and wrote descriptions in a standard format that described their functionality and how they were used, taking care to not imply any judgment about the features' reliability or usability. Features were described either as existing commercially or developed by researchers; in all cases we asked participants to suppose the privacy feature was available with the device they were reviewing. We used a within-subjects design where participants were shown three privacy features in sequence to reduce the percent of time participants spent on study setup and collect a larger number of responses with our research budget.

3.1.3 Information About Privacy Features' Limitations. To describe the features' limitations, we likewise reviewed the literature and noted limitations mentioned in prior work (e.g., ease of use, performance, security). We determined that the limitations fit into two overarching categories: functionality (how the feature works when operating as designed, including usability issues), and reliability (when the feature might not work as designed,

Table 3. Participants were first shown a description of the privacy feature, and later shown more detail about its limitations. The description for hardware mute control is shown as an example; descriptions for all features are in Appendix C

Description	Limitations
Some devices have a hardware mute control , such as a switch or button that lets you control whether the cameras and microphones are on. This switch, which cannot be accessed or controlled remotely, makes the sensors completely inoperable. When you want to disable the cameras and microphones, you have to walk up to the smart indoor camera and toggle the control, which will disconnect power to the sensors.	Functionality: If you turn off the cameras and microphones using the hardware control, then you cannot interact with this Sustios smart indoor camera. So, if you turned off the sensors because you were concerned about the smart device recording when you were on a Zoom meeting, you wouldn't be able to use the device until you walked over to it and turned the sensors back on. Reliability: The hardware mute control does not have a reliability limitation and should always work.

such as due to security and performance issues). We then wrote descriptions following a standard format. A sample description is shown in Table 3; Appendix C contains descriptions for all features.

The functionality limitation restated the feature's interaction mechanism as a limitation. For example, we mentioned that to use the activity log feature, you would have to open an app periodically and audit the log. The reliability limitation for software-backed and audio heuristics/jamming features mentioned that software bugs or vulnerabilities may prevent the feature from working reliably. For hardware-backed features, we mentioned that the control "does not have a reliability limitation and should always work," as third-party audits of Google products have verified the mute control's hardware implementation revealed no vulnerabilities that would result in the microphone being turned on [53]. For the indicator light, we mentioned that the indicator light itself will always work (e.g., if the microphone is off, show a red light), but that software bugs may cause the sensors to be turned on unexpectedly. Prior work has found vulnerabilities in some indicator light implementations [14], but newer devices, including Amazon Echo products, use hardware-backed implementations that ensure the light and sensor state are in sync; while a software bug may turn on the sensors, the light would also change state [9, 25].

We used a within-subjects design for this portion to be able to measure participants' individual changes in responses to the privacy features and their limitations by comparing individual Likert-scale and free-response responses before and after limitations were described.

3.2 Analysis Methods and Metrics

We used quantitative and qualitative methods to analyze our survey results. In this section, we first describe our analysis methodology, and then describe the approach we use for each research question.

3.2.1 Analysis Methodology. Mirroring prior work evaluating the factors that influence privacy risk perception for IoT [28], when modeling the response to a single question on a Likert scale, we used Cumulative Link Mixed Models (CLMMs), which allow modeling all levels of the ordinal response, as opposed to a logistic regression model that would require binning the responses [4]. To model changes in trust over the duration of the survey, we used logistic regressions to model whether there was an increase or decrease in agreement. All multiple-choice questions used five-point Likert scales and included a "Not sure" option. "Not sure" responses (less than 2% of all responses) were excluded from further analysis.

Additionally, we coded a sample of the qualitative responses to provide additional context for trends observed in our quantitative analysis. We sampled at least 25% of responses to each question, with at least 25 responses for each group we wished to analyze separately (e.g., participants who expressed that the hardware mute control was unreliable). In all cases, we ensured that we reached saturation of themes and ideas with a subset of our sample, then continued coding until we reviewed at least 25% of responses. As we used a single coder and only

reviewed a sample of responses, we do not report the exact number of responses that matched a particular code or sentiment. We use the following terminology to describe the proportion of responses that matched a particular code: “a few” for 0–20% of responses, “some” for 20–40%, “about half” for 40–60%, and “many” for 60–100%.

We report on statistically significant results where $p < 0.05$, and include the odds ratios. For the CLMMs, the odds ratio represents the odds of a response having higher agreement for that factor compared to the baseline. For the logistic regressions, the odds ratio corresponds to the odds of the participant’s trust increasing or decreasing (depending on the model) when additional information is provided. All models included independent variables for the device type and brand of the product and a boolean for whether the participant owned one of the products discussed in the survey. For the questions that were shown multiple times with different privacy features, we included a random effect to model per-participant variation in responses. When comparing devices, we used a Sustios smart speaker as a baseline to be able to assess changes in both brand and device, and when comparing privacy features, we used software mute control as a baseline as we expected it would rank near the middle of features with most response variables. Full statistical results are included in Appendix D.

3.2.2 Device Factors (RQ1). We created a series of CLMM models for the responses about privacy concerns with and trust in the device, before privacy features were described. The response variable was the (ordinal) response for concern or agreement. For the free-response question about trust in the product, we coded 150 responses, randomly sampling 25 responses for each combination of device type and brand.

3.2.3 Privacy Features (RQ1). We used CLMMs to model the Likert-scale (ordinal) responses about intent to use and perceived reliability. As we were interested in reactions when participants had full information about the privacy and usability impacts of each feature, we modeled the responses *after* participants saw information about the features’ limitations. To model how privacy features impacted trust, we used a logistic regression. The response variable (boolean) was whether the participant responded with *increased* agreement to the statement that they trusted the device, compared to their response earlier in the survey before the privacy features were introduced. Our hypothesis was that privacy controls would increase trust, so we binned responses where there was no change or a decrease in trust to “non-increase.”

We coded a sample of 400 responses about the privacy features’ reliability, sampling 50 responses for each privacy feature, including 25 responses where participants considered the feature reliable after limitations were described, and 25 responses where they did not, based on their responses to the corresponding multiple-choice question. We reviewed the responses about reliability both before and after the limitations were described, and assigned codes to their combined responses.

3.2.4 Limitations of Privacy Features (RQ2). We likewise used logistic regressions to model the change in responses. Our response variable was whether the response to WouldUse-Feature, Reliable-Feature, and Trust-Device indicated *decreased* agreement than the response to the corresponding question before limitations were described. In this case, our hypothesis was that agreement would decrease, so responses indicating no change in agreement were bucketed with those with an increase.

To provide additional context regarding whether users anticipated the tradeoffs, we coded a random sample of 400 responses to the question asking users if they thought the privacy feature had any limitations (before we described them), with 50 responses for each privacy feature. To better understand how individual participants’ opinions changed based on the information we provided, we coded 400 responses about whether participants would use the privacy feature by reviewing the answers before and after limitations were described and categorizing the change in their responses. We again sampled 50 responses per privacy feature, with half of those responses selected from participants who indicated they would use the privacy feature.

3.3 Study Limitations

Our study uses a convenience sample of smart-home users from Prolific, which may not be representative of the general population. However, prior work has shown that responses from crowdworker platforms are generally representative of the U.S. population age 18–49 [59], that Prolific participants are more diverse and have higher-quality responses than respondents from competing crowdworker platforms [3, 57], and that Prolific participants' responses to questions about privacy perceptions and beliefs are generalizable to the larger population [67]. We recruited participants who own at least one smart-home device as we expected they would be more familiar with the content of our survey. While our study design allows comparing responses of users who own a sensor-driven product (e.g., smart camera) with users who only own other products (e.g., smart lighting), our results may not generalize to users who do not own any such products. We measured intended behaviors and reactions after participants reviewed a limited amount of information; actual behaviors in real-world situations may differ.

We assessed only a few types of smart-home devices, brands, and sensors, as we wanted to focus on scenarios participants would likely be familiar with involving widely-available products that collect high-fidelity audio/visual data. User perceptions of other sensors, such as Lidar, temperature, and motion, where the inferences may be more sensitive than the raw data, likely are different. We also did not evaluate all possible device type/brand combinations (e.g., Amazon smart display) or measure the interaction between device type and brand, as the focus of our study was on user reactions to privacy features and not device-brand combinations. We also assumed, and indicated to participants, that hardware-backed controls would never fail; if (ostensibly) hardware-backed features in widely-available products have long-term reliability issues or failures, user trust in these features would likely be lower than what is reflected in our results.

Similar to prior work [28], we did not measure interactions between different factors in our analysis of privacy features. However, such interactions may exist—for example, participants' response to features such as the jammer wristband may vary based on the device type. Finally, we coded a sample of responses using a single coder and we use this data to provide additional context for trends revealed by the quantitative analyses; our analysis of 25%–30% of the responses may not be representative of all participants.

4 Results

We first describe participants' demographics (Section 4.1) and review how device type, brand, and ownership of similar devices impact privacy concerns and trust (Section 4.2). We then discuss reactions to privacy features after participants see the descriptions and limitations and how these features impact trust (Section 4.3). Finally, we examine how information about the limitations of privacy features impacted participants' perceptions (Section 4.4).

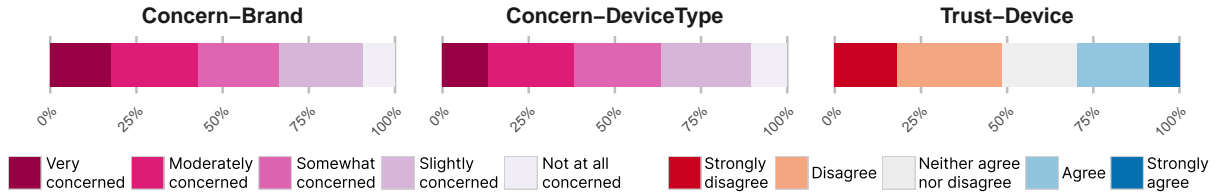
4.1 Participants

We ran our study on Prolific in April and May 2024. 500 participants completed our survey. We had four comprehension-check questions; 11 participants failed one or more of the checks, so we discarded their responses and analyzed the remaining 489 responses. Our survey took a median of 14.8 minutes. We asked participants whether we could share their anonymized responses publicly; we have published the full responses for the 452 participants who agreed (removing their Prolific IDs and any potentially-identifying information) and analysis code to replicate our statistical results.¹

The participants reflected the demographics of crowdworking platforms. We offered the survey to a gender-balanced sample of users on Prolific; 49% of participants identified as male, 48% as female, and 2% as non-binary. 46% were ages 18–34, and 53% had a bachelor's degree or higher, and 28% had education or a job in computer science or a related field. These demographics are consistent with prior studies run on crowdworking platforms [28, 59].

¹<https://osf.io/zycbq/>

Fig. 1. Responses about the devices we presented indicate a high level of distrust. Each bar corresponds to a question, with segments showing the proportion of participants selecting each response option. Table 1a shows the full question text.



We recruited participants who indicated on their Prolific profile that they owned a smart-home product (defined broadly to include devices such as smart TVs). 81% of participants indicated that they owned at least one smart-home product, with 75% owning at least one of the products discussed in our survey: 65% owned a smart speaker, 20% a smart display, and 35% a smart indoor camera. 50% owned more than one type of device.

4.2 Device Factors

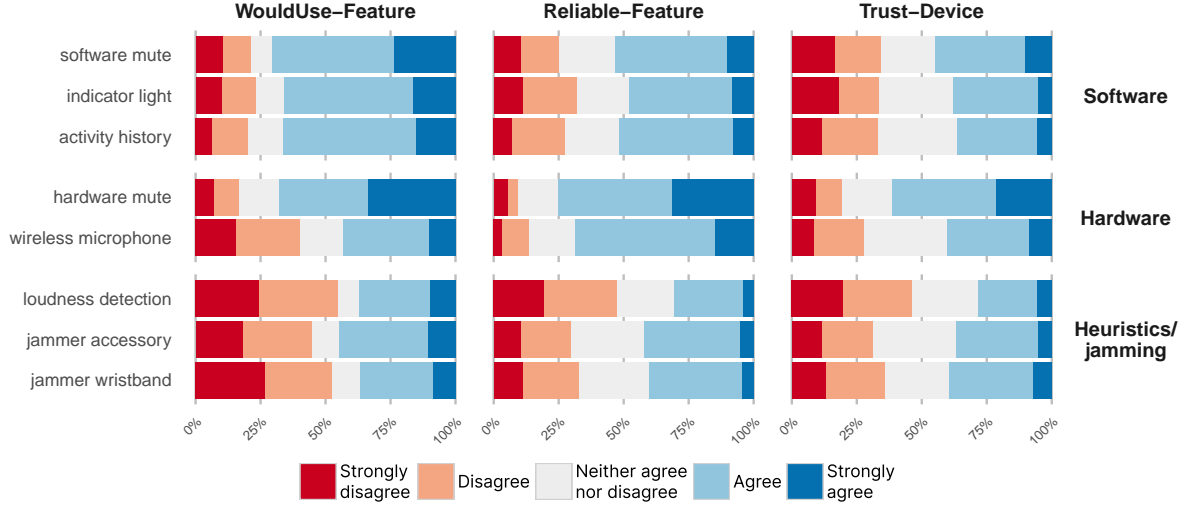
In this section, we discuss participants' privacy concerns with and trust in the devices in our survey, including how device type, brand, and ownership of similar devices impact these reactions. Overall, participants expressed distrust in how companies handle personal data. Participants were more likely to express trust in devices similar to those they owned, and expressed slightly higher trust in smart indoor cameras compared to smart speakers.

A majority of respondents were somewhat, moderately, or very concerned about how the products in our survey collect, store, and use information. About half (48%) of participants indicated that they did not trust the product they reviewed to protect their privacy. Figure 1 shows the responses to these questions. This distrust exists across all the companies we asked about in our survey. Depending on the brand, 19% to 37% of participants indicated they trusted the device, with our fictional "Sustios" brand in the middle of this range; however, none of the differences were statistically significant. The overall sentiment was negative: across all brands that we asked about, more respondents indicated distrust than indicated trust. Free-text responses that mentioned a brand name more frequently indicated that the company's reputation decreased trust rather than increased it. Some participants expressed general distrust in all companies or IoT devices, such as P238, who cited "the history of other companies' actions."

Participants who owned any smart speaker, smart display, or smart indoor camera were more likely to have lower levels of privacy concerns and higher levels of trust. In our CLMM models for responses about privacy concerns, device owners were about three times as likely to indicate a lower level of privacy concern in the device ($OR = 0.26, p < 0.001$) and in the brand ($OR = 0.33, p < 0.001$) than non-owners. Similarly, device owners were also about three times as likely to indicate increased trust in how the device collects data ($OR = 3.21, p < 0.001$), how it uses and shares information ($OR = 3.09, p < 0.001$), and the device generally ($OR = 3.58, p < 0.001$). The odds ratio (OR) represents the odds of the response indicating more trust, with values > 1 indicating that the factor increases trust compared to the baseline. Additionally, participants assessing a smart indoor camera indicated slightly higher trust that the device would only collect data when appropriate ($OR = 1.83, p = 0.030$) and protect their privacy ($OR = 1.87, p = 0.026$).

Finding 1: Many participants expressed distrust of all products, regardless of device type and brand. However, participants were more likely to indicate trust in a product if they owned a similar one.

Fig. 2. Reactions to privacy features varied by type of feature. The bars show responses for each privacy feature/question pair, with segments showing the proportion of responses with each level of agreement. Table 1b shows the full question text.



4.3 Privacy Features

In this section, we review participants' reaction to privacy features. Participants expressed higher intent to use most software- and hardware-backed features, but expressed more mixed reactions to features using audio heuristics and jamming (Figure 2). Hardware-backed features were considered the most reliable and increased trust. Unless otherwise noted, we discuss responses *after* information about the features' limitations was presented.

4.3.1 Factors Impacting Responses to All Privacy Features. Participants had a variety of reactions to the privacy features we described, with no feature eliciting uniformly positive or negative responses. Among the 1,464 total responses for all privacy features (each participant saw three features), participants said that they would use a privacy feature 53% of the time and considered the feature reliable 51% of the time. A few participants expressed a desire to be able to verify a feature's behavior, such as through "see[ing] the feature activated and work as intended" (P495), or "other sources like trusted reviews" (P243). Continuing the sentiment expressed earlier in the survey, participants were divided on whether they trusted IoT devices and their manufacturers: among the text responses that discussed these factors, about half indicated some type of distrust. However, in general, the presence of a privacy feature increased trust.

Participants who owned a smart speaker, smart display, or smart indoor camera expressed that privacy features were more reliable than non-owners did, with owners being about twice as likely to indicate a higher level of agreement than non-owners ($OR = 2.01, p < 0.001$). Participants who reviewed a smart indoor camera considered the privacy features to be less reliable than the baseline smart speaker ($OR = 0.57, p = 0.031$, or 1.75 times as likely to be in a lower level of agreement), and were less likely to indicate that the privacy features increased trust ($OR = 0.21, p = 0.034$, or 4.8 times as likely to indicate a lower level of trust). In contrast, device brand did not have a statistically significant impact on any of the responses.

4.3.2 Software-Backed Features. Participants ranked software-backed features (software mute control, indicator light, and activity history screen) relatively highly, particularly with respect to their intent to use the features, but did not consider them as reliable as hardware-backed controls. 70% of respondents agreed or strongly agreed with

the statement that they would use the software mute control, the highest among all privacy controls. Similarly, 66% and 65% of participants indicated they would use the activity history screen and indicator light, respectively. Participants considered the software-backed features to be moderately reliable and trustworthy, with about half of participants agreeing or strongly agreeing with the statement that the feature is reliable for the software mute control (53%), activity history screen (51%), and indicator light (47%).

We did not find a statistically significant difference between responses for the indicator light or activity history screen and the baseline software mute control. Even though these three features have very different interaction mechanisms, participants' responses to the three controls were similar, suggesting high intent to use, moderate perceived reliability, and moderate impact on trust in the product. In subsequent sections, we compare other classes of privacy features to the software mute control, which we use as the baseline.

4.3.3 Hardware-Backed Features. Participants considered privacy features backed by a hardware interlock (hardware mute control and wireless microphone accessory) to be the most reliable and trustworthy, though a minority of participants expressed concerns. More than two thirds of participants considered these features reliable. Compared to the baseline of a software mute control, participants considered the hardware mute control ($OR = 5.36, p < 0.001$) and wireless microphone ($OR = 2.36, p < 0.001$) more reliable. Participants also indicated that they would trust devices with these controls more. Participants were 7.82 times more likely to indicate increased trust (compared to their response before we described privacy features) in a device with a hardware mute control ($OR = 7.82, p < 0.001$), and had a higher probability of indicating increased trust with the wireless microphone ($OR = 2.68, p = 0.013$), compared to the software mute control.

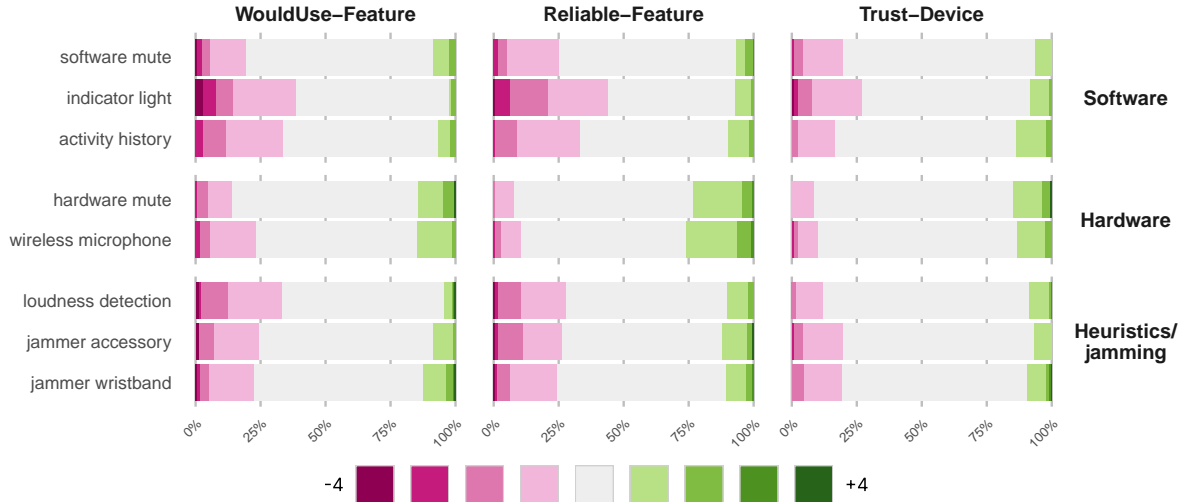
However, for this class of privacy features, responses about intent to use hardware-backed controls were similar to comparable software-backed features. We found no statistically significant difference between the reported intent to use a hardware and software mute controls; 68% and 70% of responses indicated they would use the hardware and software controls, respectively. Participants described multiple usage patterns for the hardware mute control: a few participants expressed that they would always or nearly always keep the microphone muted, while others would use the mute control when they wanted to ensure a specific conversation or activity was not recorded. Similarly, 43% of participants expressed intent to use the wireless microphone accessory, an identical proportion to the jammer accessory, which has a similar interaction mechanism.

Moreover, echoing prior work [41], a few participants did not trust that the hardware mute control works as advertised, such as P331, who said “Just because they say it will work, that doesn’t mean it will actually work.” Among the text responses we sampled from participants who did not agree the feature was reliable, about half expressed that they didn’t trust the control and a few would not trust it unless they could verify its behavior. These findings suggest that while hardware-backed features are generally considered reliable and their presence increases trust in smart-home sensor products, some people do not trust that the controls work as advertised.

4.3.4 Audio Heuristics and Jamming. Participants had more mixed reactions to privacy features that used audio heuristics or jamming (loudness detection feature, jammer accessory, and jammer wristband). 40% of participants considered the jammer devices reliable when we initially presented the features, citing reasons including that they are “simple” (P109, jammer wristband) and “sciencey” (P357, jammer wristband). However, others expressed concerns, including unreliability due to manufacturing variations (P338).

On balance, however, participants responded less positively to this class of privacy feature. For the jammer accessory and jammer wristband, they reported lower intent to use ($OR = 0.29, p < 0.001$ and $OR = 0.21, p < 0.001$, respectively), with only 43% responses indicating intent to use the jammer accessory and 36% for the jammer wristband, compared to 70% for the software mute control. Likewise, respondents considered the jamming devices less reliable ($OR = 0.60, p = 0.016$ and $OR = 0.56, p = 0.007$, respectively), with 41% and 40% agreeing or strongly agreeing the accessory and wristband, respectively were reliable, compared to 53% for the

Fig. 3. The plot shows the degree to which participants' responses changed, with pink indicating a decrease in agreement and green an increase. Participants generally indicated less intent to use privacy features and considered them less reliable after seeing information about their limitations, with more responses changing for the software-backed controls.



software mute control. As we will discuss in the next section, some of this more-critical reaction to the jammer devices is likely due to information we provided about their limitations.

Participants expressed similar reservations with the loudness detection feature, reporting lower intent to use and reliability compared to the baseline ($OR = 0.22, p < 0.001$ and $OR = 0.30, p < 0.001$, respectively), with only 37% indicating they would use it and 30% considering it reliable. They also indicated lower trust: participants were 73% less likely to express that they trusted the product more with a loudness detection feature than a software mute control ($OR = 0.27, p = 0.003$). In the text responses, some were concerned about the feature's accuracy, citing potential false positives, e.g., due to “an intense situation at your house (arguments)” (P460).

Finding 2: Participants generally express intent to use both software and hardware-backed controls; they indicate that hardware-backed controls are more reliable and engender more trust, but nevertheless indicate similar intent to use software-backed features. Participants also expressed lower intent to use and lower perceived reliability for features using audio heuristics or jamming.

4.4 Limitations of Privacy Features

For most privacy features, learning more about the limitations caused participants to indicate they were less likely to use the feature, consider it reliable, or trust the device (Figure 3). However, participants considered hardware-backed features more reliable after seeing text mentioning the lack of reliability limitations.

4.4.1 Software-Backed Features. The tradeoffs of the software-backed features were often not initially obvious to participants: many participants did not mention any limitations at all or mentioned incorrect limitations when we asked them after first presenting the feature.

After we described the limitations, participants indicated lower intent to use the *software mute control* ($OR = 0.22, p < 0.001$), considered it less reliable ($OR = 0.15, p < 0.001$), and trusted it less ($OR = 0.17, p < 0.001$). These effect sizes correspond to a 18% probability that participants would indicate lower intent to use for the software mute control, 13% probability they would consider it less reliable, and 15% probability they would trust a device with a software mute control less. In the free-text responses, there was some divergence between responses about intent to use and trust, such as a participant who said “I would still use it but I wouldn’t trust it” (P428).

The reduction in intent to use and reliability was even more pronounced for the *indicator light*. Initially, in their text responses, many participants did not volunteer any limitations. When we mentioned that it might be possible for the microphone to turn on (and the light indicating the microphone is muted to turn off) without the user noticing it, participants were more likely to indicate the feature was less reliable than the baseline ($OR = 3.05, p < 0.001$). Likewise, participants indicated they would use the indicator light less often ($OR = 2.97, p < 0.001$). Even with information about those tradeoffs, however, participants still ranked the indicator light more favorably than most controls, with 65% agreeing that they would use it, compared to 70% for the software mute control.

Similarly, for the *activity history screen*, intent to use similarly decreased compared to the baseline ($OR = 2.31, p = 0.002$), again suggesting that the limitations of software-backed features are not always initially understood, though with less strong of an effect than the indicator light. Among the sampled text responses, some indicated that the limitations were not a problem for them, such as one participant who said that “Software bugs are to be expected in any technology.” (P370).

4.4.2 Hardware-Backed Features. In contrast, information about the limitations of the hardware-backed features (*hardware mute control* and *wireless microphone*) had a smaller impact on responses. The functionality limitation for both features mentioned that you need to walk up to the device to turn the microphone on and off. Somewhat surprisingly, this description did not seem to sway participants. For the hardware mute control, some participants indicated they would be less likely to use the feature (23 responses), but almost as many indicated they would be more likely (24 out of a total of 169 responses).

The vast majority of participants considered the hardware mute control reliable, both before (69%) and after (75%) we displayed the text discussing limitations (which mentioned that the hardware control should always work). With respect to intent to use, the baseline agreement was already high (69% of participants), and on balance the information about the limitations did little to change participants’ responses. These trends were reflected in our models as well: compared to explaining the limitations of the baseline software mute control, explaining limitations of hardware devices was less likely to cause a decreased belief in reliability ($OR = 0.21, p < 0.001$) and trust ($OR = 0.38, p = 0.005$). Similarly, responses for the wireless microphone also indicated a lower decrease in reliability ($OR = 0.30, p < 0.001$) and trust ($OR = 0.43, p = 0.008$).

Trust in devices with a hardware mute control also increased (56% to 60% of participants), in contrast to all other features, where trust decreased. As before, a minority of participants expressed distrust in the hardware mute control, and our statement that the control does not have reliability limitations did not always sway their opinions. P451, for example, mentioned that “I have a feeling that the company still somehow has a way to switch it off and listen in regardless of the actual hardware switch.”

4.4.3 Audio Heuristics and Jamming. Responses to information about the limitations of non-traditional privacy features (*jammer accessory* and *jammer wristband*) were similar to other software-backed features: intent to use, reliability, and trust all decreased, but there was not a statistically significant difference in this effect from the baseline software mute control. More responses considered the *loudness detection feature* less reliable, and participants expressed concern about false positives and negatives.

All three of the audio heuristics and jamming features started with a lower baseline level of intent to use, reliability, and trust: before we displayed the limitations, only 46%, 50%, and 38% of participants said they would use the loudness detection feature, jammer accessory, and jammer wristband, respectively, compared to 72% who

would use the software mute control. There was not a statistically significant difference in the *change* in intent to use compared to the software mute control. This may suggest that participants understood these features' functionality and reliability tradeoffs to a similar degree as those of currently-available features, but did not feel that they have a use case for these features.

Participants' text responses indicated some, but not universal, understanding of the tradeoffs for these features. For the jammer accessory and wristband, about half of responses mentioned a specific limitation, with the most common being that it impacts product functionality, e.g., it "may disturb phone calls" (P238) or cause issues "if a burglar wore one" (P52). In contrast, only a few sampled responses for the software mute control mentioned a concern with the feature limiting the functionality of the device. Similarly, for the loudness detection feature, about half of participants mentioned reliability concerns, such as "[not] picking up soft speaking people" (P176) or listening all the time "if have a normally loud voice" (P141). Participants' intent to use the loudness detection feature decreased further than the baseline ($p = 0.004$).

4.4.4 Other Factors. We additionally observed that perceived reliability of the privacy features decreased more for smart displays (OR = 2.03, $p = 0.018$) and smart indoor cameras (OR = 3.28, $p < 0.001$) than for the baseline smart speaker, and more for Amazon (OR = 1.73, $p = 0.002$) products than for the fictional Sustios device. On the other hand, device ownership did not have a statistically significant impact on responses.

Finding 3: The impact of conveying limitations varied by feature. After being told of their limitations, participants expressed lower intent to use, perceived reliability, and trust for software-backed features and for audio heuristics/jamming features. Participants considered hardware-backed features more reliable, but did not indicate significantly higher intent to use those features.

5 Discussion

In this section, we summarize the results of our study, and then discuss the implications of our findings for the design of privacy features for smart-home sensor products.

5.1 Factors Impacting Trust in Smart-Home Sensor Products

Participants showed a general distrust in how companies and smart-home devices handle data. However, participants who owned a similar smart-home sensor product expressed higher trust, echoing prior work [20, 73]. Contrary to our expectations and prior work [34], device brand did not have a significant impact. Intent to use privacy features and trust were somewhat decoupled, with participants expressing similar intent to use hardware- and software-backed privacy features, but indicating more trust in products with hardware controls.

Hardware-backed features induce the most trust. Based on participants' responses to our survey, hardware mute controls increase users' trust in smart-home products the most. Many participants (60%) agreed that they would trust the product if it had a hardware mute control. Participants also ranked hardware mute controls highly in terms of reliability (highest among all controls) and intent to use (second highest). However, a minority of participants (9%) expressed that they did not trust the hardware mute control to work reliably, even after we mentioned in the survey that it should always work as intended. Prior work also observed the same sentiment [41].

Software-backed features provide balance between reliability and usability. The software mute control, indicator light, and activity history screen all ranked highly in terms of intent to use (65% to 70% of responses indicating agreement). As expected, the software-backed features were not considered as reliable as the hardware-backed ones, and these features increased trust in the smart-home products less the hardware mute control. However, the high rankings for intent to use suggest that the features provide value to users.

Research solutions fit some use cases, but not home IoT environments. While prior work studying prototype privacy features found that they increased users' perceptions of privacy [19, 49], in our study, participants were more critical of these features. It is possible that in-person use of the features helps people trust that they are working as intended (at least while the user is testing the feature), and future work should consider how these more complex features are best explained to users. Regarding the jammer devices, a few participants were concerned about more devices being jammed than intended (e.g., disabling their cell phone when they want to make a phone call). This behavior may be desired in some use cases, such as at a vacation rental where users may be concerned about hidden recording devices. However, in the home IoT context, devices that jam all microphones may be undesired; solutions targeted at a single device may be more appropriate. The wireless microphone, in contrast, has a hardware guarantee that the microphone is only on when the user expects it to be [22]. As expected, participants considered the wireless microphone to be more reliable than the jammer devices. However, responses about intent to use and trust were similar to responses for the jammer devices. With respect to the loudness detection feature [49], participants expressed concerns that the feature, intended to filter audio so a voice assistant would only be triggered with a louder-than-normal voice, would not function as expected.

5.2 Recommendations

Similar to prior work [22], we believe that the goals of privacy features should be to make products more *trustworthy* and *trusted*. Participants' responses to our study suggest that privacy features and data handling practices need to be verifiable to meet these goals of being trustworthy and trusted. We discuss recommendations for device manufacturers, researchers, and regulators based on our findings.

Improve transparency around privacy features' limitations (device manufacturers, researchers, regulators, consumer-protection and advocacy groups). Our results suggest that the reliability tradeoffs of software-backed features are not obvious to users. Disclosures by device manufacturers [33] increase transparency, but our work shows that *feature-specific* details, not just general information about the device, matter. Additionally, the disclosures may have nuances that are not obvious to users; for example, documentation for the Amazon Echo mentions that when the microphones are turned off, the device's circuitry ensures that the red indicator lights turn on [9]. However, no claim is made about any hardware guarantees for the microphone mute button, and third-party audits suggest that the device can turn the microphones back on at any time [25]. Future work should explore how to communicate the tradeoffs of privacy features to ensure users do not make incorrect assumptions about what guarantees a feature provides or underestimate potential risks such as software bugs or hardware failures. While device manufacturers could add more detail to their companion mobile apps, other mechanisms such as privacy dashboards or ambient lighting [68] may be more effective.

Provide verifiable privacy assurances (device manufacturers). Our study found that many people have a general distrust in how all companies handle data, a finding also observed in prior work [66]. A portion of participants were not satisfied by any of the privacy features we presented: 47 of 489 did not consider any of the features they saw to be reliable; among those participants, 85% indicated at the beginning of the survey that they did not trust the device to protect their privacy. Future work should explore how to build trust in smart-home products through other means, such as independent verification of privacy practices and clearer communication about how data is used and shared. Other work has found that consumers are willing to pay for increased privacy and security in IoT devices, and that a "nutrition label" can more clearly communicate privacy practices [29]. Ongoing efforts such as the U.S. Cyber Trust Mark [18] may encourage more third-party privacy and security audits, which have the potential to address some of the concerns participants raised in our survey.

Evaluate real-world usage of privacy features (researchers). In our study, 68% of participants expressed agreement or strong agreement with the statement that they would use a hardware mute control. Participants who own one

of the products discussed in our survey, which commonly have hardware mute controls, expressed equally-strong agreement (66%). In contrast, in a previous study, only two out of seventeen interviewed smart-speaker users used the mute button, and neither user fully trusted it [41]. In another study, only 5% of smart speaker users reported using the mute button [45]. As these prior studies were conducted in 2018–2019, it is possible that familiarity with hardware mute controls and other privacy features has increased over time. It is additionally possible that even with our study design intended to have people consider the usability tradeoffs of the hardware mute control, participants still overstate their intent to use the control. As such, we recommend that future work explore strategies to measure actual usage of privacy features in smart-home sensor products.

Improve security guarantees of software-backed privacy features (device manufacturers, researchers). Participants indicated high intent to use software-backed features, but expressed more trust in devices with a hardware mute control. Future work should explore how to improve the security guarantees of software-backed features to shrink this gap in trust. Numerous projects have explored ways to secure the sensor data processing pipeline without changing the way users interact with the device. For example, De Vaere et al. prototyped an auditable voice-activation pipeline using trusted execution environments [69], Warden et al. proposed running ML models on sensor data on isolated microcontrollers [71], and Jin et al. placed more data processing on a local smart-home hub [39]. Device manufacturers and researchers should explore how these techniques and others can be deployed more widely. Building and deploying these systems may make smart-home sensor products more *trustworthy* as they provide stronger security and privacy guarantees. However, if users do not understand how these systems work, or make incorrect conclusions about the systems being more secure or insecure, they may not be *trusted* by users. In our study, participants mentioned that they found privacy features that were “simple” to be reliable. Novel systems using trusted execution environments or dedicated microcontrollers introduce new technical complexity, and further work is needed to understand how to explain these systems to users.

Improve verifiability of add-on privacy features (researchers). The research community has developed add-on privacy features and accessories, such as jammer devices, that can provide control without requiring the underlying smart-home product to be modified. However, features that use audio heuristics or jamming have some inherent unreliability. A few participants in our study mentioned that they would want to test the feature themselves or have a trusted third party verify its reliability. Future work should consider how these features can offer users ways to confirm they are functioning correctly (e.g., accurately blocking nearby microphones).

Verify software-backed features work as expected (regulators, consumer-protection and advocacy groups). More participants indicated they would use software-backed features than indicated the features were reliable, suggesting that consumers are willing to use the features that best fit their use case even if they are not fully reliable or trusted. Consumers may be unwilling or unable to evaluate security risks and assess the software quality of products they purchase. Regulators and consumer protection and advocacy groups should consider how to ensure that implementations of software features meet best practices for ensuring reliability, such as by developing minimum security standards, IoT nutrition labels, or testing infrastructure [13, 18, 26].

6 Conclusion

Microphones and cameras in smart home products introduce privacy risks. To address those risks and users’ concerns, device manufacturers and researchers have developed a variety of privacy features to provide additional transparency and control. In this paper, we evaluated the extent to which these privacy features contribute to trust in smart-home sensor products. Participants’ responses to our survey suggest that people trust products which offer verifiable privacy assurances; in particular, hardware mute controls are the most effective at increasing trust in smart-home sensor products. However, participants also expressed distrust in how all companies handle data, and a portion of participants were not satisfied by any of the privacy features we evaluated. These results

suggest that there is no single factor that makes a smart home product trustworthy, and that manufacturers and researchers should explore how to make privacy features more secure and verifiable.

Acknowledgments

We gratefully acknowledge support from Craig Newmark Philanthropies and the National Science Foundation (award SaTC-1801472). We thank Jenny Tang, McKenna McCall, and Eric Zeng for their assistance discussing study methodology, piloting the survey, and reviewing a draft of this paper.

References

- [1] Noura Abdi, Xiao Zhan, Kopo M. Ramokapane, and Jose Such. 2021. Privacy Norms for Smart Home Personal Assistants. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [2] Desiree Abrokwa, Shruti Das, Omer Akgul, and Michelle L. Mazurek. 2021. Comparing Security and Privacy Attitudes Among U.S. Users of Different Smartphone and Smart-Speaker Platforms. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS)*.
- [3] Troy Adams, Yuanxia Li, and Hao Liu. 2020. A Replication of Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research – Sometimes Preferable to Student Groups. *AIS Transactions on Replication Research* (2020).
- [4] Alan Agresti. 2012. *Analysis of Ordinal Categorical Data* (2nd ed. ed.). Wiley.
- [5] Imtiaz Ahmad, Taslima Akter, Zachary Buher, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2022. Tangible Privacy for Smart Voice Assistants: Bystanders’ Perceptions of Physical Device Controls. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022).
- [6] Imtiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2020. Tangible Privacy: Towards User-Centric Sensor Designs for Bystander Privacy. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020).
- [7] Shima Ahmed, Amrita Roy Chowdhury, Kassem Fawaz, and Parmesh Ramanathan. 2020. Preech: A System for Privacy-Preserving Speech Transcription. In *29th USENIX Security Symposium*.
- [8] Amazon. [n. d.]. *Alexa Privacy*. <https://www.amazon.com/Alexa-Privacy-Hub/b?node=19149155011>
- [9] Amazon. [n. d.]. *How Alexa works: Microphone on off button*. <https://www.amazon.com/b/?node=23608613011>
- [10] Tawfiq Ammari, Jofish Kaye, Janice Y. Tsai, and Frank Bentley. 2019. Music, Search, and IoT: How People (Really) Use Voice Assistants. *ACM Transactions on Computer-Human Interaction* (2019).
- [11] Apple. 2019. *Improving Siri’s privacy protections*. <https://www.apple.com/newsroom/2019/08/improving-siris-privacy-protections/>
- [12] Apple. 2024. *HomePod Privacy and Security*. <https://support.apple.com/guide/homepod/privacy-and-security-apd99ee29027/homepod>
- [13] Mehrdad Bahrini, Nima Zargham, Alexander Wolff, Dennis-Kenji Kipker, Karsten Sohr, and Rainer Malaka. 2022. It’s Long and Complicated! Enhancing One-Page Privacy Policies in Smart Home Applications. In *Nordic Human-Computer Interaction Conference*.
- [14] Matthew Brocker and Stephen Checkoway. 2014. iSeeYou: Disabling the MacBook Webcam Indicator LED. In *23rd USENIX Security Symposium*.
- [15] Sara Cannizzaro, Rob Procter, Sinong Ma, and Carsten Maple. 2020. Trust in the smart home: Findings from a nationally representative survey in the UK. *PLOS ONE* (2020).
- [16] Narae Cha, Auk Kim, Cheul Young Park, Soowon Kang, Mingyu Park, Jae-Gil Lee, Sangsu Lee, and Uichin Lee. 2020. Hello There! Is Now a Good Time to Talk? Opportune Moments for Proactive Interactions with Smart Speakers. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020).
- [17] Varun Chandrasekaran, Suman Banerjee, Bilge Mutlu, and Kassem Fawaz. 2021. PowerCut and Obfuscator: An Exploration of the Design Space for Privacy-Preserving Interventions for Smart Speakers. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS)*.
- [18] Claire C. Chen, Dillon Shu, Hamsini Ravishankar, Xinran Li, Yuvraj Agarwal, and Lorrie Faith Cranor. 2024. Is a Trustmark and QR Code Enough? The Effect of IoT Security and Privacy Label Information Complexity on Consumer Comprehension and Behavior. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
- [19] Yuxin Chen, Huiying Li, Shan-Yuan Teng, Steven Nagels, Zhijiang Li, Pedro Lopes, Ben Y. Zhao, and Haitao Zheng. 2020. Wearable Microphone Jamming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [20] Camille Cobb, Sruti Bhagavatula, Kalil Anderson Garrett, Alison Hoffman, Varun Rao, and Lujo Bauer. 2021. “I would have to evaluate their objections”: Privacy tensions between smart home device owners and incidental users. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (2021).
- [21] Matt Day, Giles Turner, and Natalia Drozdak. 2019. Amazon Workers Are Listening to What You Tell Alexa. *Bloomberg* (2019). <https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alex-a-global-team-reviews-audio>
- [22] Youngwook Do, Nivedita Arora, Ali Mirzazadeh, Injoo Moon, Eryue Xu, Zhihan Zhang, Gregory D. Abowd, and Sauvik Das. 2023. Powering for Privacy: Improving User Trust in Smart Speaker Microphones with Intentional Powering and Perceptible Assurance. In *32nd USENIX Security Symposium*.

- [23] Youngwook Do, Jung Wook Park, Yuxi Wu, Avinandan Basu, Dingtian Zhang, Gregory D. Abowd, and Sauvik Das. 2021. Smart Webcam Cover: Exploring the Design of an Intelligent Webcam Cover to Improve Usability and Trust. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2021).
- [24] Serge Egelman, Raghudeep Kannavara, and Richard Chow. 2015. Is This Thing On? Crowdsourcing Privacy Indicators for Ubiquitous Sensing Platforms. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*.
- [25] electronupdate. 2021. *Amazon Echo Flex: Microphone Mute, Real or Fake?* <https://electronupdate.blogspot.com/2021/01/amazon-echo-flex-microphone-mute-real.html>
- [26] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the Experts: What Should Be on an IoT Privacy and Security Label?. In *2020 IEEE Symposium on Security and Privacy (SP)*.
- [27] Pardis Emami-Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS)*.
- [28] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2021. Which Privacy and Security Attributes Most Impact Consumers' Risk Perception and Willingness to Purchase IoT Devices?. In *2021 IEEE Symposium on Security and Privacy (SP)*.
- [29] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of IoT Devices?. In *32nd USENIX Security Symposium*.
- [30] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [31] Davide Ferraris, Daniel Bastos, Carmen Fernandez-Gago, and Fadi El-Moussa. 2021. A trust model for popular smart home devices. *International Journal of Information Security* (2021).
- [32] Jonas Foehr and Claas Christian Germelmann. 2020. Alexa, Can I Trust You? Exploring Consumer Paths to Trust in Smart Voice-Interaction Technologies. *Journal of the Association for Consumer Research* (2020).
- [33] Google. [n. d.]. *Google Nest Security & Privacy Commitments*. <https://safety.google/nest/>
- [34] Wentao Guo, Jason Walter, and Michelle L. Mazurek. 2023. The Role of Professional Product Reviewers in Evaluating Security and Privacy. In *32nd USENIX Security Symposium*.
- [35] Rajibul Hasan, Riad Shams, and Mizan Rahman. 2021. Consumer trust and perceived risk for voice-controlled artificial intelligence: The case of Siri. *Journal of Business Research* (2021).
- [36] Alex Hern. 2019. Apple contractors 'regularly hear confidential details' on Siri recordings. *The Guardian* (2019). <https://www.theguardian.com/technology/2019/jul/26/apple-contractors-regularly-hear-confidential-details-on-siri-recordings>
- [37] Kashmir Hill. 2020. Activate This 'Bracelet of Silence,' and Alexa Can't Eavesdrop. *The New York Times* (2020). <https://www.nytimes.com/2020/02/14/technology/alexa-jamming-bracelet-privacy-armor.html>
- [38] Haojian Jin, Boyuan Guo, Rituparna Roychoudhury, Yaxing Yao, Swarun Kumar, Yuvraj Agarwal, and Jason I. Hong. 2022. Exploring the Needs of Users for Supporting Privacy-Protective Behaviors in Smart Homes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.
- [39] Haojian Jin, Gram Liu, David Hwang, Swarun Kumar, Yuvraj Agarwal, and Jason I. Hong. 2022. Peekaboo: A Hub-Based Approach to Enable Transparency in Data Processing within Smart Homes. In *2022 IEEE Symposium on Security and Privacy (SP)*.
- [40] Bjørn Karmann. 2018. *Project Alias*. https://bjoernkarmann.dk/project/project_alias
- [41] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018).
- [42] Sunok Lee, Minji Cho, and Sangsu Lee. 2020. What If Conversational Agents Became Invisible? Comparing Users' Mental Models According to Physical Entity of AI Speaker. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020).
- [43] Yuqi Liu, Yan Gan, Yao Song, and Jing Liu. 2021. What Influences the Perceived Trust of a Voice-Enabled Smart Home System: An Empirical Study. *Sensors* (2021).
- [44] Nathan Malkin. 2019. *Privacy controls for always-listening devices*. Ph.D. Dissertation. University of California, Berkeley.
- [45] Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. 2019, 4 (2019).
- [46] Nathan Malkin, David Wagner, and Serge Egelman. 2022. Can Humans Detect Malicious Always-Listening Assistants? A Framework for Crowdsourcing Test Drives. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022).
- [47] Graeme McLean and Kofi Osei-Frimpong. 2019. Hey Alexa ... examine the variables influencing the use of artificial intelligent in-home voice assistants. *Computers in Human Behavior* (2019).
- [48] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys that Listen: A Study of Parents, Children, and Internet-Connected Toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*.
- [49] Abraham Mhaidli, Manikandan Kandadai Venkatesh, Yixin Zou, and Florian Schaub. 2020. Listen Only When Spoken To: Interpersonal Communication Cues as Smart Speaker Privacy Controls. *Proceedings on Privacy Enhancing Technologies* 2020 (2020). Issue 2.

- [50] Oliver Michler, Reinhold Decker, and Christian Stummer. 2020. To trust or not to trust smart consumer products: a literature review of trust-building factors. *Management Review Quarterly* (2020).
- [51] MSCHF. [n. d.]. *Alexagate*. <https://alexagate.com>
- [52] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2024. SoK: Data Privacy in Virtual Reality. *Proceedings on Privacy Enhancing Technologies* 2024, 1 (2024).
- [53] NCC Group ioXt Validation Lab. 2021. Nest Mini ioXt Device Assessment. <https://storage.googleapis.com/support-kms-prod/aUdYcY7JigJgD1vsMuMZujo4QqahlGthh6xi>
- [54] Jared Newman. 2020. You can now buy an Amazon Echo add-on that stops Alexa from listening. *Fast Company* (2020). <https://www.fastcompany.com/90532150/you-can-now-buy-an-amazon-echo-add-on-that-stops-alexa-from-listening/>
- [55] Paranoid. [n. d.]. *Paranoid Home Device*. <https://paranoid.com/products>
- [56] Prasoon Patidar, Mayank Goel, and Yuvraj Agarwal. 2023. VAX: Using Existing Video and Audio-based Activity Recognition Models to Bootstrap Privacy-Sensitive Sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2023).
- [57] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* (2017).
- [58] Rebecca S. Portnoff, Linda N. Lee, Serge Egelman, Pratyush Mishra, Derek Leung, and David Wagner. 2015. Somebody’s Watching Me? Assessing the Effectiveness of Webcam Indicator Lights. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*.
- [59] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *2019 IEEE Symposium on Security and Privacy (SP)*.
- [60] Nirupam Roy, Haitham Hassanieh, and Romit Roy Choudhury. 2017. BackDoor: Making Microphones Hear Inaudible Sounds. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*.
- [61] Ali Shahin Shamsabadi, Brij Mohan Lal Srivastava, Aurélien Bellet, Nathalie Vauquier, Emmanuel Vincent, Mohamed Maouche, Marc Tommasi, and Nicolas Papernot. 2023. Differentially Private Speaker Anonymization. *Proceedings on Privacy Enhancing Technologies* 2023, 1 (2023).
- [62] Rahul Anand Sharma, Elahe Soltanaghaei, Anthony Rowe, and Vyas Sekar. 2022. Lumos: Identifying and Localizing Diverse Hidden IoT Devices in an Unfamiliar Environment. In *31st USENIX Security Symposium*.
- [63] Akash Deep Singh, Brian Wang, Luis Garcia, Xiang Chen, and Mani Srivastava. 2024. Understanding factors behind IoT privacy – A user’s perspective on RF sensors. <http://arxiv.org/abs/2401.08037>
- [64] Ke Sun, Chen Chen, and Xinyu Zhang. 2020. “Alexa, Stop Spying on Me!”: speech privacy protection against voice assistants. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*.
- [65] Madiha Tabassum, Tomasz Kosiński, Alisa Frik, Nathan Malkin, Primal Wijesekera, Serge Egelman, and Heather Richter Lipford. 2019. Investigating Users’ Preferences and Expectations for Always-Listening Voice Assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (2019).
- [66] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. 2019. “I don’t own the data”: End User Perceptions of Smart Home Device Data Practices and Risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS)*.
- [67] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS)*.
- [68] Parth Kirankumar Thakkar, Shijing He, Shiyu Xu, Danny Yuxing Huang, and Yaxing Yao. 2022. “It would probably turn into a social faux-pas”: Users’ and Bystanders’ Preferences of Privacy Awareness Mechanisms in Smart Homes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.
- [69] Piet De Vaere and Adrian Perrig. 2023. Hey Kimya, Is My Smart Speaker Spying on Me? Taking Control of Sensor Privacy Through Isolation and Amnesia. In *32nd USENIX Security Symposium*.
- [70] Tim Verheyden, Denny Baert, Lente Van Hee, and Ruben Van Den Heuvel. 2019. Google employees are eavesdropping, even in your living room, VRT NWS has discovered. *VRT NWS* (2019). <https://www.vrt.be/vrtnws/en/2019/07/10/google-employees-are-eavesdropping-even-in-flemish-living-rooms/>
- [71] Pete Warden, Matthew Stewart, Brian Plancher, Colby Banbury, Shvetank Prakash, Emma Chen, Zain Asgar, Sachin Katti, and Vijay Janapa Reddi. 2022. Machine Learning Sensors. <http://arxiv.org/abs/2206.03266>
- [72] Shilin Xiao, Xiaoyu Ji, Chen Yan, Zhicong Zheng, and Wenyan Xu. 2023. MicPro: Microphone-based Voice Privacy Protection. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*.
- [73] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018).
- [74] Ruochen Zhou, Xiaoyu Ji, Chen Yan, Yi-Chao Chen, Wenyan Xu, and Chaohao Li. 2023. DeHiREC: Detecting Hidden Voice Recorders via ADC Electromagnetic Radiation. In *2023 IEEE Symposium on Security and Privacy (SP)*.

A Survey Instrument

Part 1: Device privacy concerns and trust

Suppose [brand] is introducing a new [device].

The product supports the following features: *[Description of one randomly-selected device type is shown. See Appendix B for possible descriptions.]*

Based on your existing knowledge of smart home products, [brand], [device] products, and the information provided about this [brand] [device], please answer the following questions.

Which of these choices best describes how you feel about how [device] products in general collect, store and use information? *(Concern-DeviceType: 5-point scale, Not at all concerned to Very concerned)*

Which of these choices best describes how you feel about how [brand] products collect, store and use information? *(Concern-Brand: 5-point scale, Not at all concerned to Very concerned)*

I would trust this [brand] [device] to only collect data (e.g. audio recordings) when appropriate. *(Trust-Collect: 5-point scale, Strongly agree to Strongly disagree)*

I would trust this [brand] [device] to use and share data it collects appropriately. *(Trust-Use: 5-point scale, Strongly agree to Strongly disagree)*

I would trust this [brand] [device] to protect my privacy. *(Trust-Device: 5-point scale, Strongly agree to Strongly disagree)*

What considerations contribute to how much you would trust this [brand] [device]? *(Free response)*

Part 2: Privacy features

[Repeated three times with a different randomly-selected privacy feature.]

To address privacy concerns with smart devices, a number of technologies have been developed to give you more control and transparency when the [device]'s sensors are on.

We will show you a number of technologies that you could use to better understand and control about how and when this [brand] [device] records you.

[Description of one randomly-selected privacy feature is shown. See Appendix C for possible descriptions.]

Which technology is this section about? *(Multiple choice, comprehension check question)*

Think about how you could use this [feature] with a new [brand] [device] product.

I would use the [feature] with a [brand] [device]. *(WouldUse-Feature: 5-point scale, Strongly agree to Strongly disagree)*

Please explain why you would or would not use the [feature]. *(Free response)*

I trust that the [feature] would work reliably. *(Reliable-Feature: 5-point scale, Strongly agree to Strongly disagree)*

Please explain why you do or do not trust that the [feature] would work reliably. *(Free response)*

Given I could use the [feature] with this product, I would trust that this [brand] [device] protects my privacy. *(Trust-Device: 5-point scale, Strongly agree to Strongly disagree)*

Do you think the [feature] has any limitations (e.g. or you would lose some of the [device]’s functionality by using it)? *(Free response)*

Privacy-protective technologies often have limitations that impact the usability of [device] products. In this section, we will discuss some of the limitations with this [feature].

[Description the limitations of the previously-shown privacy feature is shown. See Appendix C for possible descriptions.]

Knowing these limitations, I would use the [feature] with a [brand] [device]. *(WouldUse-Feature: 5-point scale, Strongly agree to Strongly disagree)*

Please explain why you would or would not use the [feature], knowing these limitations. *(Free response)*

Knowing these limitations, I trust that the [feature] described would work reliably. *(Reliable-Feature: 5-point scale, Strongly agree to Strongly disagree)*

Please explain why you do or do not trust that the [feature] would work reliably, knowing these limitations. *(Free response)*

Given the limitations of using the [feature], I would trust that this [brand] [device] protects my privacy. *(Trust-Device: 5-point scale, Strongly agree to Strongly disagree)*

Part 3: Device ownership and demographics

Which types of smart home products do you own?

- Smart speaker (e.g. Amazon Echo)
- Smart display (e.g. Google Nest Hub)
- Smart camera (e.g. Wyze Cam)
- Other smart home products (e.g. smart light bulb)
- I do not own any smart home products

Which brands of smart home products do you own?

- Amazon
- Google
- Apple
- Other

What is your age?

- 18–24 years
- 25–34 years
- 35–44 years
- 45–54 years
- 55–64 years
- 65 years or older
- Prefer not to say

How would you describe your gender? *(Man, Woman, Non-binary, Prefer to self-describe (text entry), Prefer not to say)*

What is the highest level of education you have completed?

- Some high school
- High school diploma or equivalent
- Some college
- Trade, technical, or vocational training
- Associate degree
- Bachelor's degree
- Master's degree
- Professional degree
- Doctoral degree
- Prefer not to say

Which of the following best describes your educational background or job field?

- I have an education in, or work in, the field of computer science, engineering, or IT.
- I do not have an education in, or work in, the field of computer science, engineering, or IT.
- Prefer not to say

After this research is completed, we plan on publishing anonymized responses to this survey (including multiple choice and free-response questions) from participants who agree to have their data published. Do you agree to have your anonymized responses to this survey posted publicly? Your compensation will not be impacted by your response to this question.

- I agree to have my anonymized responses posted publicly
- Do not post my responses publicly

B Device Descriptions

The device descriptions were written to be similar across device types while mentioning key features that use the cameras and microphones on the products.

Participants saw a bulleted list with all the items for a single device type, regardless of the brand and device they saw. Where possible, we used the same features and similar wording across device types.

Table 4. Device descriptions

Smart speaker	Smart display	Smart indoor camera
Voice assistant: the device listens for when you say a wake word, letting you ask the device questions about the weather, automate your home, and more.	Smart assistant: the device listens for when you say a wake word, letting you ask the device questions about the weather, automate your home, and more.	—
—	—	Video recording: records video when a person or pet is detected.
—	Person identification: the device looks to see if a recognized user is nearby, and shows your calendar if it detects you.	Person identification: if the person or pet in front of the camera has been seen before, the camera can send a notification to its owner saying who was detected.
Sound detection: the device listens for smoke alarms or breaking glass, letting you know if there is an issue in your home.	Sound detection: the device listens for smoke alarms or breaking glass, letting you know if there is an issue in your home.	Sound detection: the device listens for smoke alarms or breaking glass, letting you know if there is an issue in your home.
Smart home integration: you can control all popular brands of smart home accessories, so you can ask the device to turn on the lights or change the thermostat.	Smart home integration: you can control all popular brands of smart home accessories, so you can ask the device to turn on the lights or change the thermostat.	Smart home integration: The camera can be viewed and controlled with all major smart home platforms.

C Privacy Feature Descriptions

Participants were presented with the description of the privacy feature, answered questions about the feature, then presented with the description of limitations and answered additional questions. This portion of the study used a within-subjects design where participants were randomly assigned to review three of the eight features.

Table 5. Privacy feature descriptions

Feature	Description	Limitations
hardware mute control	Some devices have a hardware mute control , such as a switch or button that lets you control whether the [cameras and] microphones are on. This switch, which cannot be accessed or controlled remotely, makes the sensors completely inoperable. When you want to disable the [cameras and] microphones, you have to walk up to the [device] and toggle the control, which will disconnect power to the sensors. Suppose that this [brand] [device] has a hardware mute control .	Functionality: If you turn off the [cameras and] microphones using the hardware control, then you cannot interact with this [brand] [device]. So, if you turned off the sensors because you were concerned about the smart device recording when you were on a Zoom meeting, you wouldn't be able to use the device until you walked over to it and turned the sensors back on. Reliability: The hardware mute control does not have a reliability limitation and should always work.
software mute control	Some devices have a software mute control that you can operate using a mobile app. Using this app, you can easily turn the [cameras and] microphones on and off. You also can create home screen widgets and shortcuts to make it easy to control the sensors, and set automations to have them turn on and off on a schedule. When you want to turn off the [cameras and] microphones on your [device], you can open the app and turn off the sensors, which will tell the device to stop listening. Suppose that this [brand] [device] has a software mute control .	Functionality: If you want to turn the [cameras and] microphones back on, you have to tap a button on your phone or set an automation to turn it on. Reliability: The software mute control should almost always work, but software bugs or vulnerabilities may prevent the control from working reliably.
indicator light	Some devices have indicator light that ensures you can tell whether the [cameras and] microphones are on by looking at the device. The device is wired so that power is provided either to the dedicated red light on the device or to the device's [cameras and] microphones, but not both at the same time. That means it's physically impossible for the sensors and the red light to be on at the same time. Suppose that this [brand] [device] has an indicator light .	Functionality: The light accurately reflects whether the sensors are on, but if the sensors are turned on by another person or due to a software bug, you might not notice the light is turned off and the [cameras and] microphones are on again. Reliability: The indicator light should always work, but software bugs or vulnerabilities could cause the sensors to be turned on unexpectedly.

Table 5. Privacy feature descriptions (*continued*)

Feature	Description	Limitations
jammer accessory	<p>Researchers have created a jammer accessory that can be placed near a [device] and controlled using an app on your phone. You can use this jammer to disable the microphones on nearby smart devices, so you don't have to trust that the privacy controls on the [device] are working as designed.</p> <p>When the device is turned on, it blocks the microphones on the nearby smart device using inaudible ultrasonic noise. When you want to disable nearby microphones, you can use the app on your phone to turn on the jammer, which will prevent the microphones from hearing you.</p> <p>Suppose that you can use a jammer accessory with this [brand] [device].</p>	<p>Functionality: If you want to unblock the microphones, you have to tap a button on your phone or set an automation to disable the jammer.</p> <p>Reliability: When operated as designed, the jammer accessory should almost always work, but software bugs or vulnerabilities may prevent the control from working reliably.</p>
jammer wristband	<p>Researchers have created a jammer wristband that you can wear. You can wear this wristband to prevent all nearby microphones from hearing you, so you don't have to that the privacy controls on the [device] are working as designed.</p> <p>When the wristband is turned on, it blocks the microphones on nearby smart devices using inaudible ultrasonic noise.</p> <p>Suppose that you can use a jammer wristband with this [brand] [device].</p>	<p>Functionality: If you want to unblock the microphones, you have to tap a button on your phone or set an automation to disable the jammer.</p> <p>Reliability: When operated as designed, the jammer wristband should almost always work, but software bugs or vulnerabilities may prevent the control from working reliably.</p>
wireless microphone	<p>Researchers have created a wireless microphone that you can turn on and off by opening and closing a hinge, like a very small laptop computer. When you interact with the microphone by opening the hinge, it is wirelessly powered and turns on.</p> <p>When you want to have a [device]'s microphones enabled, you can go up to the wireless microphone, open it, and talk to the product.</p> <p>Suppose that you can use a wireless microphone with this [brand] [device].</p>	<p>Functionality: If you want to turn on the microphone, you have to walk up to it and open the hinge, which could be inconvenient.</p> <p>Reliability: The wireless microphone does not have a reliability limitation and should always work.</p>

Table 5. Privacy feature descriptions (*continued*)

Feature	Description	Limitations
activity history screen	<p>Some devices have a activity history screen that you can view using a mobile app. Using this app, you can easily turn the [cameras and] microphones on and off. You also can create home screen widgets and shortcuts to make it easy to control the sensors, and set automations to have them turn on and off on a schedule.</p> <p>When you want to view when the [cameras and] microphones on your [device] are used, you can open the app and look at the activity history and review recordings of the sensor data.</p> <p>Suppose that this [brand] [device] has a activity history screen.</p>	<p>Functionality: If you want to audit the [device]'s use of its [cameras and] microphones, you have to review the activity log and consider when you used the device to determine whether the product is using its sensors as expected.</p> <p>Reliability: The activity history screen should almost always work, but software bugs or vulnerabilities may prevent the app from working reliably.</p>
loudness detection feature	<p>Researchers have proposed a loudness detection feature that would be built into [device] products. With this feature enabled, the device will only think you are interacting with it if the microphones hear you talking louder than normal speech, so normal conversations won't be recorded.</p> <p>Suppose that this [brand] [device] has a loudness detection feature.</p>	<p>Functionality: If you want to interact with the [device] you have to speak louder than normal for it to hear you.</p> <p>Reliability: The loudness detection feature should almost always work, but software bugs or vulnerabilities may cause the device to record when you are speaking quietly.</p>

D Detailed Regression Results

D.1 Device Factors

For the device-specific response variables (Concern-Device, Concern-Brand, Trust-Collect, Trust-Use, and Trust-General), we used a Cumulative Link Mixed Model to determine correlation with increased concerns and trust. An estimate > 0 indicates that the factor is associated with higher agreement. The odds ratios indicate the (multiplicative) change in the odds of agreement, i.e., an odds ratio of 0.259 for a factor indicates that the odds of agreement are 25.9% of the odds of agreement with the baseline.

Table 6. Regression results for Concern-Device (“Which of these choices best describes how you feel about how [device] products in general collect, store and use information?”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: 1 2	—	−3.123	—	0.283	−11.026	< 0.001
Threshold: 2 3	—	−1.478	—	0.255	−5.800	< 0.001
Threshold: 3 4	—	−0.373	—	0.248	−1.505	0.132
Threshold: 4 5	—	1.145	—	0.258	4.444	< 0.001
Device: Sustios display	Sustios speaker	−0.054	0.948	0.276	−0.195	0.845
Device: Sustios camera	Sustios speaker	0.088	1.092	0.272	0.325	0.745
Device: Amazon speaker	Sustios speaker	0.252	1.287	0.279	0.904	0.366
Device: Apple speaker	Sustios speaker	0.365	1.441	0.284	1.287	0.198
Device: Google speaker	Sustios speaker	0.215	1.240	0.291	0.741	0.459
Device owner	Non-owner	−1.348	0.260	0.197	−6.831	< 0.001

Table 7. Regression results for Concern-Brand (“Which of these choices best describes how you feel about how [brand] products collect, store and use information?”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: 1 2	—	−3.014	—	0.282	−10.680	< 0.001
Threshold: 2 3	—	−1.358	—	0.249	−5.447	< 0.001
Threshold: 3 4	—	−0.338	—	0.243	−1.394	0.163
Threshold: 4 5	—	0.980	—	0.249	3.936	< 0.001
Device: Sustios display	Sustios speaker	−0.079	0.924	0.277	−0.284	0.776
Device: Sustios camera	Sustios speaker	0.008	1.008	0.273	0.031	0.976
Device: Amazon speaker	Sustios speaker	0.433	1.542	0.283	1.528	0.126
Device: Apple speaker	Sustios speaker	0.274	1.316	0.283	0.969	0.333
Device: Google speaker	Sustios speaker	0.484	1.623	0.286	1.695	0.090
Device owner	Non-owner	−1.094	0.335	0.193	−5.659	< 0.001

Table 8. Regression results for Trust-Collect (“I would trust this [brand] [device] to only collect data (e.g. audio recordings) when appropriate.”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: -2 -1	—	-0.780	—	0.254	-3.069	0.002
Threshold: -1 0	—	0.906	—	0.253	3.581	< 0.001
Threshold: 0 1	—	1.746	—	0.261	6.696	< 0.001
Threshold: 1 2	—	3.768	—	0.308	12.238	< 0.001
Device: Sustios display	Sustios speaker	0.434	1.543	0.275	1.580	0.114
Device: Sustios camera	Sustios speaker	0.607	1.835	0.280	2.169	0.030
Device: Amazon speaker	Sustios speaker	-0.083	0.921	0.279	-0.297	0.767
Device: Apple speaker	Sustios speaker	0.427	1.532	0.282	1.514	0.130
Device: Google speaker	Sustios speaker	-0.133	0.876	0.298	-0.445	0.656
Device owner	Non-owner	1.168	3.215	0.199	5.857	< 0.001

Table 9. Regression results for Trust-Use (“I would trust this [brand] [device] to use and share data it collects appropriately.”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: -2 -1	—	-0.798	—	0.249	-3.207	0.001
Threshold: -1 0	—	0.937	—	0.249	3.770	< 0.001
Threshold: 0 1	—	1.806	—	0.258	7.008	< 0.001
Threshold: 1 2	—	3.816	—	0.310	12.297	< 0.001
Device: Sustios display	Sustios speaker	0.492	1.636	0.272	1.811	0.070
Device: Sustios camera	Sustios speaker	0.461	1.586	0.277	1.669	0.095
Device: Amazon speaker	Sustios speaker	-0.380	0.684	0.283	-1.346	0.178
Device: Apple speaker	Sustios speaker	0.276	1.318	0.280	0.986	0.324
Device: Google speaker	Sustios speaker	-0.182	0.834	0.297	-0.613	0.540
Device owner	Non-owner	1.127	3.087	0.199	5.673	< 0.001

Table 10. Regression results for Trust-General (“I would trust this [brand] [device] to protect my privacy”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: -2 -1	—	-0.543	—	0.248	-2.187	0.029
Threshold: -1 0	—	1.050	—	0.253	4.147	< 0.001
Threshold: 0 1	—	2.051	—	0.265	7.736	< 0.001
Threshold: 1 2	—	3.561	—	0.298	11.961	< 0.001
Device: Sustios display	Sustios speaker	0.269	1.309	0.269	1.002	0.317
Device: Sustios camera	Sustios speaker	0.624	1.866	0.281	2.223	0.026
Device: Amazon speaker	Sustios speaker	-0.293	0.746	0.284	-1.032	0.302
Device: Apple speaker	Sustios speaker	0.388	1.473	0.284	1.364	0.172
Device: Google speaker	Sustios speaker	-0.282	0.754	0.293	-0.962	0.336
Device owner	Non-owner	1.274	3.575	0.203	6.288	< 0.001

D.2 Privacy Features

For WouldUse-Feature and Reliable-Feature, we used a Cumulative Link Mixed Model to determine correlation with higher agreement with the respective reaction statements. An estimate > 0 indicates that the factor is associated with higher agreement with the reaction statement. The odds ratios indicate the (multiplicative) change in the odds of agreement.

For Trust-Device, we used a logistic regression to determine correlation with higher agreement with the respective reaction statements. Our response variable (boolean) was whether agreement with the reaction statement about trust in the smart-home product increased compared to the participant's response to a similar question before any privacy features were described. An estimate > 0 indicates that the factor is associated with an increase in agreement with the reaction statement.

Table 11. Regression results for WouldUse-Feature (“I would use the [feature] with a [brand] [device].”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: -2 -1	—	-2.695	—	0.258	-10.426	< 0.001
Threshold: -1 0	—	-1.356	—	0.248	-5.473	< 0.001
Threshold: 0 1	—	-0.728	—	0.245	-2.977	0.003
Threshold: 1 2	—	1.536	—	0.249	6.165	< 0.001
Feature: indicator light	Software mute	-0.310	0.733	0.208	-1.492	0.136
Feature: activity history screen	Software mute	-0.246	0.782	0.207	-1.187	0.235
Feature: hardware mute control	Software mute	0.338	1.402	0.215	1.568	0.117
Feature: wireless microphone	Software mute	-1.020	0.361	0.205	-4.968	< 0.001
Feature: loudness detection	Software mute	-1.505	0.222	0.212	-7.097	< 0.001
Feature: jammer accessory	Software mute	-1.240	0.290	0.211	-5.885	< 0.001
Feature: jammer wristband	Software mute	-1.574	0.207	0.214	-7.353	< 0.001
Device: Sustios display	Sustios speaker	-0.030	0.971	0.220	-0.135	0.893
Device: Sustios camera	Sustios speaker	-0.240	0.787	0.219	-1.094	0.274
Device: Amazon speaker	Sustios speaker	0.082	1.085	0.225	0.363	0.717
Device: Apple speaker	Sustios speaker	-0.069	0.933	0.223	-0.309	0.757
Device: Google speaker	Sustios speaker	-0.050	0.951	0.228	-0.221	0.825
Trust-General: Agree	Non-agree	0.567	1.763	0.148	3.833	< 0.001
Device owner	Non-owner	0.072	1.075	0.154	0.470	0.638

Table 12. Regression results for Reliable-Feature (“I trust that the [feature] would work reliably.”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Threshold: -2 -1	—	-2.488	—	0.293	-8.490	< 0.001
Threshold: -1 0	—	-0.750	—	0.280	-2.682	0.007
Threshold: 0 1	—	0.635	—	0.278	2.285	0.022
Threshold: 1 2	—	3.769	—	0.303	12.437	< 0.001
Feature: indicator light	Software mute	-0.330	0.719	0.218	-1.512	0.130
Feature: activity history screen	Software mute	-0.118	0.888	0.219	-0.540	0.589
Feature: hardware mute control	Software mute	1.679	5.362	0.234	7.173	< 0.001
Feature: wireless microphone	Software mute	0.858	2.358	0.219	3.915	< 0.001
Feature: loudness detection	Software mute	-1.204	0.300	0.220	-5.473	< 0.001
Feature: jammer accessory	Software mute	-0.515	0.597	0.215	-2.401	0.016
Feature: jammer wristband	Software mute	-0.584	0.558	0.218	-2.682	0.007
Device: Sustios display	Sustios speaker	-0.257	0.773	0.264	-0.973	0.331
Device: Sustios camera	Sustios speaker	-0.565	0.568	0.262	-2.154	0.031
Device: Amazon speaker	Sustios speaker	-0.422	0.656	0.269	-1.568	0.117
Device: Apple speaker	Sustios speaker	0.065	1.067	0.267	0.242	0.808
Device: Google speaker	Sustios speaker	-0.233	0.792	0.276	-0.845	0.398
Trust-General: Agree	Non-agree	1.571	4.810	0.182	8.607	< 0.001
Device owner	Non-owner	0.696	2.005	0.185	3.761	< 0.001

Table 13. Regression results for Trust-Device (“Given I could use the [feature] with this product, I would trust that this [brand] [device] protects my privacy.”).

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Intercept	—	-1.450	0.235	0.709	-2.046	0.041
Feature: indicator light	Software mute	0.330	1.391	0.405	0.815	0.415
Feature: activity history screen	Software mute	0.213	1.237	0.400	0.532	0.595
Feature: hardware mute control	Software mute	2.057	7.822	0.430	4.785	< 0.001
Feature: wireless microphone	Software mute	0.987	2.683	0.396	2.492	0.013
Feature: loudness detection	Software mute	-1.314	0.269	0.445	-2.952	0.003
Feature: jammer accessory	Software mute	0.295	1.343	0.395	0.748	0.455
Feature: jammer wristband	Software mute	0.270	1.310	0.409	0.660	0.509
Device: Sustios display	Sustios speaker	-0.565	0.568	0.747	-0.756	0.450
Device: Sustios camera	Sustios speaker	-1.582	0.206	0.746	-2.122	0.034
Device: Amazon speaker	Sustios speaker	0.979	2.662	0.748	1.309	0.191
Device: Apple speaker	Sustios speaker	-0.272	0.762	0.749	-0.363	0.717
Device: Google speaker	Sustios speaker	-0.542	0.582	0.776	-0.698	0.485
Device owner	Non-owner	-0.607	0.545	0.504	-1.205	0.228

D.3 Limitations of Privacy Features

To model the impact of information about limitations, we used logistic regressions. Our response variable (boolean) was whether the level of trust in the smart-home product, considering a given privacy feature and its limitations, indicated a *lower* level of trust compared to the participant's response when we initially described the privacy feature. An estimate > 0 indicates that the factor is associated with a decrease in trust. The odds ratios indicate the (multiplicative) change in the odds of a decrease, i.e., an odds ratio of 0.259 for a factor indicates that the odds of a decrease are 25.9% of the odds of a decrease with the baseline.

Table 14. Regression results for change in WouldUse-Feature after limitations are shown.

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Intercept	—	−1.524	0.218	0.301	−5.056	< 0.001
Feature: indicator light	Software mute	1.089	2.973	0.270	4.035	< 0.001
Feature: activity history screen	Software mute	0.836	2.306	0.273	3.063	0.002
Feature: hardware mute control	Software mute	−0.399	0.671	0.313	−1.274	0.203
Feature: wireless microphone	Software mute	0.233	1.262	0.276	0.842	0.400
Feature: loudness detection	Software mute	0.774	2.169	0.269	2.879	0.004
Feature: jammer accessory	Software mute	0.346	1.413	0.279	1.240	0.215
Feature: jammer wristband	Software mute	0.223	1.250	0.284	0.786	0.432
Device: Sustios display	Sustios speaker	0.148	1.160	0.250	0.593	0.553
Device: Sustios camera	Sustios speaker	0.023	1.024	0.251	0.093	0.926
Device: Amazon speaker	Sustios speaker	−0.331	0.718	0.264	−1.252	0.210
Device: Apple speaker	Sustios speaker	−0.257	0.773	0.263	−0.980	0.327
Device: Google speaker	Sustios speaker	0.279	1.322	0.258	1.081	0.280
Trust-General: Agree	Non-agree	−0.121	0.886	0.170	−0.711	0.477
Device owner	Non-owner	−0.042	0.958	0.175	−0.242	0.809

Table 15. Regression results for change in Reliable-Feature after limitations are shown.

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Intercept	—	−1.913	0.148	0.340	−5.630	< 0.001
Feature: indicator light	Software mute	1.116	3.052	0.272	4.095	< 0.001
Feature: activity history screen	Software mute	0.516	1.675	0.274	1.881	0.060
Feature: hardware mute control	Software mute	−1.574	0.207	0.375	−4.197	< 0.001
Feature: wireless microphone	Software mute	−1.199	0.302	0.323	−3.713	< 0.001
Feature: loudness detection	Software mute	0.177	1.193	0.274	0.645	0.519
Feature: jammer accessory	Software mute	0.120	1.128	0.279	0.432	0.666
Feature: jammer wristband	Software mute	0.001	1.001	0.283	0.002	0.999
Device: Sustios display	Sustios speaker	0.709	2.031	0.299	2.366	0.018
Device: Sustios camera	Sustios speaker	1.188	3.280	0.294	4.035	< 0.001
Device: Amazon speaker	Sustios speaker	0.952	2.592	0.300	3.169	0.002
Device: Apple speaker	Sustios speaker	0.161	1.175	0.313	0.514	0.608
Device: Google speaker	Sustios speaker	0.547	1.728	0.311	1.762	0.078
Trust-General: Agree	Non-agree	−0.224	0.799	0.192	−1.168	0.243
Device owner	Non-owner	−0.017	0.983	0.197	−0.088	0.930

Table 16. Regression results for change in Trust-Device after limitations are shown.

Factor	Baseline	β	Odds Ratio	Std. Error	z	p
Intercept	—	−1.769	0.171	0.316	−5.589	< 0.001
Feature: indicator light	Software mute	0.441	1.555	0.263	1.678	0.093
Feature: activity history screen	Software mute	−0.195	0.823	0.287	−0.680	0.497
Feature: hardware mute control	Software mute	−0.979	0.376	0.345	−2.837	0.005
Feature: wireless microphone	Software mute	−0.833	0.435	0.315	−2.646	0.008
Feature: loudness detection	Software mute	−0.595	0.551	0.304	−1.962	0.050
Feature: jammer accessory	Software mute	−0.069	0.933	0.278	−0.250	0.803
Feature: jammer wristband	Software mute	−0.042	0.959	0.279	−0.151	0.880
Device: Sustios display	Sustios speaker	0.224	1.252	0.266	0.845	0.398
Device: Sustios camera	Sustios speaker	0.467	1.596	0.260	1.801	0.072
Device: Amazon speaker	Sustios speaker	0.191	1.210	0.270	0.707	0.480
Device: Apple speaker	Sustios speaker	−0.189	0.828	0.287	−0.660	0.509
Device: Google speaker	Sustios speaker	0.057	1.058	0.282	0.201	0.841
Trust-General: Agree	Non-agree	−0.191	0.826	0.177	−1.080	0.280
Device owner	Non-owner	0.259	1.296	0.190	1.363	0.173