

Characterizing Everyday Misuse of Smart Home Devices

Phoebe Moh*, Pubali Datta†, Noel Warford*, Adam Bates†, Nathan Malkin*, Michelle L. Mazurek*

*University of Maryland

†University of Illinois Urbana-Champaign

Abstract—Exploration of Internet of Things (IoT) security often focuses on threats posed by external and technically-skilled attackers. While it is important to understand these most extreme cases, it is equally important to understand the most likely risks of harm posed by smart device ownership. In this paper, we explore how smart devices are misused — used without permission in a manner that causes harm — by device owners’ everyday associates such as friends, family, and romantic partners. In a preliminary characterization survey ($n = 100$), we broadly capture the kinds of unauthorized use and misuse incidents participants have experienced or engaged in. Then, in a prevalence survey ($n = 483$), we assess the prevalence of these incidents in a demographically-representative population. Our findings show that unauthorized use of smart devices is widespread (experienced by 43% of participants), and that misuse is also common (experienced by at least 19% of participants). However, highly individual factors determine whether these unauthorized use events constitute misuse. Through a focus on everyday abuses, this work sheds light on the most prevalent security and privacy threats faced by smart-home owners today.

1. Introduction

While home IoT devices benefit their users in many ways, the rapid adoption of IoT technology has also exposed new avenues for potential security and privacy violations. By their nature, these devices operate in intimate spaces, where they can collect highly sensitive information and/or control highly personal devices and environmental settings.

While researchers have raised security and privacy concerns regarding a variety of aspects of the Internet of Things, studies frequently focus on “traditional” threat actors that are external and possess great technical skill. For example, prior work has highlighted the exploitability of highly complex home automation rules [1] that do not appear to exist in practice [2, 3]. Smart speakers are also a frequent target, with reports that voice commands can be injected by operating an amplitude-modulated laser from outside the victim’s window [4], smuggling an ultrasonic speaker into the device’s vicinity [5], or embedding adversarial samples into YouTube videos [6]. Such work is important and necessary, serving to outline the limits of the attack surface of smart devices.

However, technically-demanding attacks from remote adversaries may not be the most likely of threats to a smart home. IoT devices are typically always on and can

be accessed by anyone in physical proximity, meaning that associates of the device owner such as family members, housemates, or visitors can opportunistically access these devices without permission [7–9]. While some unauthorized accesses may be more or less innocuous — for example, using a smart speaker to check the weather — users and researchers have also provided real and hypothetical examples of serious violations, including revealing sensitive information or changing device behavior in inappropriate or even threatening ways [9–12]. The seriousness of these violations often depends on expectation and circumstance, as described by the theory of contextual integrity [13].

While it is possible to imagine many types of violations, it remains unclear what kinds of interpersonal IoT *misuse* — which we define here as device use that violates social norms or the owner’s expectations of device use, leading to discomfort or harm — people are currently experiencing, how frequently misuse occurs, and how severe violations tend to be. Rather than the technical examples described above, this misuse might include non-technical violations like changing a homeowner’s smart thermostat settings without their knowledge or making a purchase on someone else’s smart speaker without permission while visiting their home. Understanding the current state of interpersonal IoT misuse can provide important context for researchers and IoT developers exploring new ways to limit access and prevent violations.

In this work, therefore, we attempt to characterize the type and frequency of interpersonal IoT misuse currently being experienced by users in the wild, starting from the following research questions:

RQ1: What kinds of misuse incidents do IoT device owners experience in their physical space? What devices do these incidents occur on, and who is doing this?

RQ2: What factors contribute to whether an unauthorized use incident becomes misuse, according to IoT device owners and users?

RQ3: How common are these IoT misuse incidents?

We explore these research questions through two surveys. We first conducted a primarily open-ended characterization survey ($n = 100$) to better understand the types of IoT misuse incidents that users experience in the physical world. From this survey, we distilled 10 categories of misuse types. Then, we used these categorizations to inform the design of a prevalence survey ($n = 483$), designed to evaluate the prevalence of these misuse events in a demographically repre-

sentative sample of the U.S. population. In both surveys, we considered the perspectives of participants who experienced misuse of their own devices and participants who engaged in misuse on another person's device.

Using these two surveys, we make the following contributions:

- 1) We characterize the types of IoT misuse incidents that are of concern to users in their day-to-day lives. These misuse incidents span a broad range of severity, from accidental exposure of viewing history to purposeful long-term spying.
- 2) We identify factors that contribute to whether an incident of unauthorized use is perceived as misuse, which include owner/user relationship, intent, information sensitivity, user mental models, and severity of consequences resulting from unauthorized use. These factors are highly individual.
- 3) We find that 43% of our prevalence survey participants report experiencing unauthorized use of one of their own home IoT devices within the five-year period before the survey was conducted, while 19% of our prevalence survey participants report having experienced misuse on one of their devices.

2. Related work

We discuss related work in three categories: remote threats related to home IoT devices (e.g., data breaches at manufacturers); local, physical threats posed by or to secondary or incidental users of home IoT devices; and misuse in non-IoT contexts.

Threats from device manufacturers and remote attackers A number of researchers have identified user concerns about home IoT devices collecting and sharing sensitive data remotely, either through intended collection (e.g., by device manufacturers) or by remote attackers. In general, people report less comfort with data collection by IoT devices in private settings, such as the home, than in public settings [14]. In the same study, Naeini et al. found that 29% of participants did not want to share IoT data with anyone due to perceived risks such as identity theft.

Personal voice assistants and smart speakers have specifically received significant attention as a source of user concern. Several studies report that users worry about whether and how manufacturers keep the data they collect safe and the potential impact of data breaches [10, 15–21]. Further, people disapprove of manufacturers repurposing data, such as using logs of voice commands to enable targeted advertising [22]. Indeed, privacy concerns like these can deter adoption [23].

In addition to threats from manufacturers, sophisticated remote adversaries can leverage attacks against IoT devices through botnets [24–26], ultrasonic carriers [21, 27], jammers [28], and leakage of cryptographic keys through various side-channels [29–31]. These sophisticated attacks, while damaging, are generally of lower concern to users, who view themselves as unlikely to be victims of targeted attacks, justifying continued IoT usage in spite of concerns [10, 25].

Privacy concerns related to remote adversaries are potentially important, but require very different design considerations than local, interpersonal IoT misuses; as such, we consider them out of scope for our work.

IoT interactions and threats in shared spaces Home IoT users face inherent security and privacy threats — whether intentional or unintentional — through day-to-day use and coexistence with family, roommates, friends, and visitors [25]. Device owners have expressed a variety of concerns over possible misuse of their devices by these secondary and *incidental users* [7, 9, 25, 32].

Incidental and secondary users themselves also have privacy and security concerns; these users are exposed to smart devices in the home but crucially do not control them. As such, the incidental users may be recorded, or otherwise have data about them collected, without their explicit permission or even their knowledge. Tensions can arise between device owners and incidental users when navigating conflicting needs, preferences, and values related to privacy and data collection [32–35]. Similar tensions were observed between guests and hosts of Airbnb rentals, where guests expressed concerns about being monitored by IoT devices included in their rental, whereas hosts desired to collect information to protect their property and ensure that house rules were being followed [36]. Further, some incidental users, such as nannies, are not in a position to express privacy concerns or preferences without economic or social risk [37].

While incidental and secondary users have the potential both to misuse devices and to suffer from misuse by device owners, prior research focusing on non-adversarial, multi-user smart homes has observed that social norms are a powerful influence in inhibiting device misuse, which decreases the need for strict access controls on these devices [7, 8, 35].

In this work, we further characterize the current state of home-IoT misuse, capturing violations that occur despite social norms and highlighting which types of hypothetical concerns our participants have actually experienced.

Misuse in non-IoT contexts Our work builds on reporting about misuse of physical access to non-IoT devices, such as mobile phones and personal computers. In particular, research has focused on *snooping*, in which an attacker takes advantage of physical presence to examine information stored on someone else's device without permission.

Researchers have found snooping on smartphones to be widespread and commonly committed by social insiders, such as friends, romantic partners, family, or others familiar with the owner [38–41]. Marques et al. evaluated the prevalence of snooping in a large population sample and estimated that 31% of survey participants had looked through someone else's smartphone without permission in the past year [38].

Even when owners give others permission to use their smartphones, unease around misuse (beyond the intended purpose of sharing the device) still persists. Smartphone owners are concerned that guest users may obtain unauthorized access to private data, delete data without permission, or

otherwise carelessly make unwanted changes, even when a phone is shared with permission [42].

Snooping also occurs on personal computers, sometimes as a means to monitor romantic partners or children [43, 44]. Broader forms of misuse, such as using a device to send out threatening messages, are also documented in shared computers found in libraries [45].

Like smartphones and computers, IoT devices are also potentially vulnerable to snooping and misuse by those with physical access to the device. Further, IoT devices are often placed intentionally to be used by multiple people in the same physical space (in contrast to phones, which in the U.S. are typically primarily single-user devices [46]). As such, we extend this prior work to focus on the types and prevalence of misuse currently occurring in home IoT devices.

3. Methods

In order to better understand the characteristics and prevalence of misuse in IoT devices, we conducted two online surveys. Survey 1, the open-ended characterization survey ($n = 100$), was used to obtain a wide variety of responses on the kinds of incidents that participants have experienced or engaged in, as well as what these participants considered to be acceptable or unacceptable in regard to their own smart devices. We used the results of this characterization survey to inform the design of Survey 2, a closed-response, larger-scale prevalence survey ($n = 483$), to assess how often the different misuse scenarios observed in the characterization survey occur.

In both surveys, we defined smart home devices as “internet-connected objects in your home, which can include lights, thermostats, smart assistants, and refrigerators. Often-times, these devices sense events in the home and change their behavior in response (for example, a doorbell camera that sends out an alert when it detects movement).” We specifically excluded computers, laptops, tablets, cell phones, and cell phone assistants from this definition.

3.1. Survey 1: the characterization survey

The characterization survey consisted of four sections:

- 1) **Instructions:** Participants were briefed about the study, presented with a consent form and definitions, and then asked about the smart devices in their home.
- 2) **Experiences:** Each participant was asked about whether they (a) noticed any unexpected changes or behaviors in one of their devices after someone used it, (b) had one of their devices used in a way that they did not expect, (c) were snooped on through one of their devices, (d) used someone else’s device while the owner was not watching, or (e) used someone else’s device to snoop. We deliberately defined these questions broadly so as to capture a wide range of experiences in this initial survey.
- a) **Drill-down:** If the participant had experienced or engaged in at least one of the topics of interest above,

we asked follow-up questions on one of the reported topics, selected uniformly at random. The follow-up questions asked for contextual information — like device type and location — in order to get a clear picture of the participant’s experience.

- b) **Secondhand stories:** If the participant had no experiences of interest, we asked if they had heard of an acquaintance experiencing something similar. If so, we asked follow-up questions about the secondhand experience.
- 3) **Privacy expectations of smart devices:** We asked participants to describe the kinds of actions they were comfortable and uncomfortable with being performed on their own and others’ devices.
- 4) **Demographics:** The survey concluded with demographic questions, which included income level, education level, and IT/CS background.

Before deployment, we tested this survey for understandability using five think-aloud interviews. We revised survey wording for clarity between interviews to address any issues discovered.

3.2. Survey 2: the prevalence survey

The prevalence survey consisted of three sections:

- 1) **Instructions:** Participants were briefed about the purpose of the study, presented with a consent form, given definitions, and then asked about what kinds of smart devices they owned.
- 2) **Experiences:** Using a matrix-style question, we presented participants with the 10 unauthorized use scenarios outlined in section 4.1, which were based on the incidents observed in the characterization survey. For each scenario, we asked participants to report if, within the last five years, they had experienced something similar on one of their own devices, engaged in something similar to another person’s device, or neither. Scenarios were presented in a random order to minimize ordering effects. One scenario served as an attention check (“Please check ‘I have done something like this...’ for this row”).
- a) **Permissions:** For each scenario the participant reported experiencing or engaging in, we asked what kind of device the scenario occurred on (presenting the eight device categories shown in Figure 1) and whether the participant gave or received (a) explicit permission, (b) implicit or assumed permission, or (c) no permission at all.
 - i) **Drill-down:** We selected at random one scenario/device combination where the participant reported an experience involving implicit or no permission for follow-up questions. Randomization weights were dynamically adjusted during survey deployment in order to capture a wider range of devices and scenarios.
 - ii) **Secondhand stories:** If the participant did not have any scenarios of interest to explore in the

“Drill-down” section, we asked if they had heard of anything like the scenarios presented in the “Experiences” section happening to or engaging in by someone they knew. We then asked similar follow-up questions about this secondhand event.

3) **Demographics:** The survey concluded with demographic questions, which included income level, education level, and IT/CS background.

Before deployment, we tested this survey for understandability with three participants using think-aloud interviews and with 30 participants through online recruitment. We made minor changes to the survey presentation as well as wording for question text and multiple-choice options before deployment.

The full text for each survey can be found in this project’s OSF repository.¹

3.3. Recruitment

Participants for both surveys were recruited through the Prolific platform and were required to reside within the U.S., be at least 18 years old, and self-report fluency in English. Because we were interested in responses from both those who have experienced misuse on their own devices and those who have engaged in misuse of another person’s device, survey participants were not required to own a smart home device.

The characterization survey took an average of 12.4 minutes to complete. Participants were paid \$5 each, which is well above the U.S. minimum hourly wage and Prolific’s suggested rates. Data collection took place in October and November 2021, and we asked participants to recall events from the past three years. While our characterization survey sample was not demographically representative, our sample was gender-balanced on gender through Prolific’s “Balanced sample” feature, which has been shown to generalize well [47].

The prevalence survey took an average of 10.6 minutes to complete, and participants were paid \$2.50 each. Data collection took place in July 2022, and we asked participants to recall events from the past five years. Participants for the prevalence survey were recruited from a demographically representative sample of the U.S. population, based on U.S. census data. The demographic breakdowns of the participants for both surveys are shown in Table 1.

We used responses to open-ended questions to validate that participants’ answers were on-topic and employed one attention check question in each survey. Responses were discarded if they provided off-topic free-response answers or if they both failed the attention check and provided low-quality answers to open-ended questions. No responses were dropped in the characterization survey, and 17 responses were dropped in the prevalence survey.

Both studies were approved by the University of Illinois Urbana-Champaign’s institutional review board (IRB).

1. https://osf.io/76daz/?view_only=ce9338830c454eba8af54908cf3be60

		S1	S2
Gender	Female	51	242
	Male	48	229
	Nonbinary	1	8
Age	18-29	57	106
	30-39	27	90
	40-49	9	72
	50-59	6	105
	60+	1	107
Number of smart devices in the home	0	7	64
	1-10	84	347
	>10	5	72
Own at least one...	Smart TV	74	357
	Smart media player	N/A	306
	Smart speaker	69	295
	Smart home management	42	195
	Smart camera	36	201
	Smart security	19	78
	Standalone smart appliance	12	73
Annual household income	Other	5	7
	<\$50k	42	180
	\$50k - \$100k	33	158
Education	>\$100k	22	131
	Have not completed high school	1	2
	High school or equivalent	42	132
	Bachelor or associate	44	254
CS background	Advanced degree	12	88
	Yes	13	79
Security background	No	82	395
	Yes	12	56
No	No	85	417

TABLE 1. PARTICIPANT DEMOGRAPHICS FOR BOTH SURVEYS.
EXCLUDES “NO ANSWER” AND “PREFER NOT TO SAY” OPTIONS.

3.4. Analysis

For the purpose of our analysis, we define unauthorized use as someone using a device without explicit permission of the owner. For the characterization survey, we determined this qualitatively through a combination of multiple choice questions about permission and supervised access with free-response descriptions of the event. For the prevalence survey, we directly asked participants whether each incident they described featured explicit, implicit, or no permission.

We define misuse as unauthorized use that violates norms or expectations, leading to discomfort or disapproval from the owner. We categorized incidents as misuse using the following criteria:

- **If the participant reported an incident where their device was used without explicit permission**, the incident was categorized as misuse if they either:
 - Reported being “Somewhat uncomfortable” or “Extremely uncomfortable” on the Likert scale question about their comfort with the incident, *or*
 - Expressed a negative sentiment about the incident, such as frustration or anger, while describing the event in open-ended questions.

- **If the participant reported an incident where they used someone else's device without explicit permission**, the incident was categorized as misuse if they stated that the device owner was either uncomfortable or had negative feelings about the incident.

These criteria were used for both the characterization survey and the prevalence survey.

For open-ended answers to the characterization survey, we performed qualitative analysis to draw out common themes [48]. Two researchers collaboratively and inductively coded 10% of the responses to develop initial codebooks for the 16 open-ended questions present in the surveys.² They then independently applied the codebooks to an additional 10% of responses at a time until strong reliability was reached at three rounds, with average Cohen's kappa of 0.84 ("strong agreement") [49]. Afterwards, the two coders divided up and independently coded the remaining responses. Very few new codes were added during the independent coding phase, confirming that conceptual saturation had been reached.

In order to identify the factors that contribute to comfort (on a five-point Likert scale) with unauthorized use, we performed ordinal logistic regression on participant responses. We compared models with a range of covariates and selected a final model based on minimum Akaike Information Criterion (AIC) [50] Further details of the analysis can be found in Section 4.3.

3.5. Limitations

Our study has limitations common to survey studies. As with any online survey, we were unable to ask our participants follow-up questions, so some responses may not fully capture nuances of participants' experiences and reactions to IoT misuse. Additionally, we were unable to obtain perspectives of the same misuse incident from both sides (victim and perpetrator). For instance, when a participant reported that they experienced snooping in the past, we were unable to solicit a response from the alleged snooper. Victims and perpetrators may not agree on what constitutes misuse.

Due to social desirability [51], participants may not disclose engaging in misuse. Thus, the observed proportions of participants that self-reported engaging in misuse should be treated as a lower bound.

A survey of this size cannot reliably capture all low-probability events, and the uneven distribution of devices in our sample may also limit generalizability of our results. For instance, Smart TVs are highly represented in our dataset by virtue of being more common than other types of devices, such as smart security devices, which may have more severe security and privacy risks. Our surveys did capture some severe events (Section 4.2) that we would expect to be low-probability. Nonetheless, we are likely missing other very severe situations which are similarly rare.

2. Full codebooks for each open-ended question can be found at this project's OSF repository, https://osf.io/76daz/?view_only=ce9338830c454ebaa8af54908cf3be60.

Recruitment was done through the online crowdsourcing platform Prolific. Though populations of online crowdsourcing platforms like Prolific are, in general, more technologically-savvy than average, they nonetheless provide reasonable sample populations [47, 52, 53]. In addition, Prolific samples have been shown to generalize well on questions relating to prior experiences and privacy and security perceptions [47].

Finally, we only recruited participants from the U.S. While this limits generalizability, it allows us to focus on one cultural context. Similar studies in other regions of the world would be beneficial to understand how IoT misuse varies across different cultures.

Characterization survey limitations We asked participants to recall events from a relatively large timeframe (the last three years) to compensate for the COVID-19 pandemic limiting opportunities for interacting with smart devices in other people's homes. However, this may have impacted recall accuracy. We asked participants to recall only one incident in detail, selected at random, in order to keep the survey length manageable. As a result, we did not exhaustively capture all IoT misuse experiences; however, this was sufficient to characterize a broad range of misuse types, as evidenced by reaching qualitative saturation [54].

Prevalence survey limitations The second survey used a multiple-choice matrix to ask about misuse incidents, relying on recognition instead of recall to improve our ability to estimate prevalence. However, as with the characterization survey, we selected only one incident to ask about in detail, using dynamic weighting to achieve broad coverage of misuse incidents while keeping survey length reasonable. As such, the detailed drill-down results provide insights about many facets of misuse incidents, but cannot be used to estimate prevalence of specific types of misuse events.

4. Results

We first characterize both unauthorized use and misuse, and then we discuss how and why participants distinguish the two. Quotes by participants from the characterization and prevalence surveys are denoted with C#X and P#X, respectively.

4.1. Characterizing unauthorized use

As noted in Section 3.4, we define unauthorized use as use without explicit permission. Respondents in the characterization survey recounted a variety of unauthorized use incidents spanning a wide range of severity, from simple pranks ("I asked Alexa to add something silly in the shopping cart for the person's Amazon account," C#99) to long-term spying ("My ex was able to extract sensitive personal information on me based upon the conversations I was having privately when she wasn't around," C#65). We used these responses, together with hypothetical scenarios respondents provided in the "Privacy expectations" section, to inductively

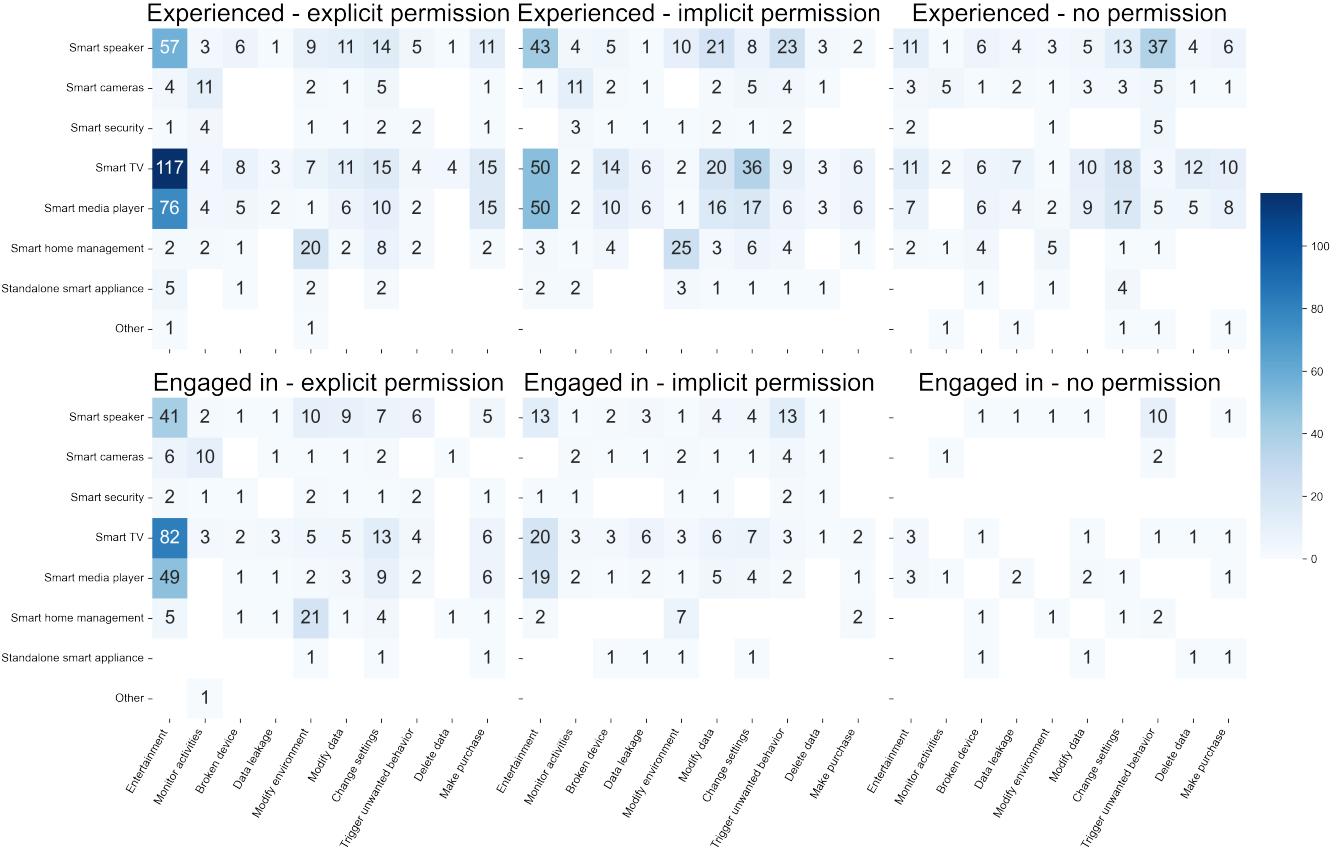


Figure 1. Number of incidents reported among our participants, organized by permission level, device type, and scenario (prevalence survey). Some device-behavior combinations may seem counterintuitive, like purchases on a smart camera, but may relate to renewed subscriptions or similar behaviors. Smart media here refers to technology like Roku or Chromecast, which provide a wide variety of media to a TV. In comparison, smart speakers are audio-only devices, like Amazon Alexa or Google Home. Smart cameras refer to video-only devices, where smart security refers to physical security, like smart locks or smart garage door openers. In-depth descriptions of scenarios can be found in Section 4.1

develop 10 categories of unauthorized use. We then used these scenario categories to solicit participants' experiences in the prevalence survey.

The scenario categories presented to participants are:

- Entertainment:** Another person used one of your smart devices to access entertainment. This can include accessing movies, music, games, or other media.
- Monitor activities:** Another person monitored your activities via one of your smart devices. This can include accessing your smart camera feed or audio recordings.
- Broken device:** One of your smart devices stopped working properly after another person used it. This can include physical damage to the device or software malfunctions.
- Data leakage:** Another person learned private information through one of your smart devices. Examples of this information can include browsing history, payment information, or account credentials.
- Modify environment:** Another person used one of your smart devices to change the physical environment of your home. This can include changing home temperature or lighting.

- Modify data:** Another person modified data on one of your smart devices. This can include adding songs to a playlist on a smart speaker or reorganizing photos on a smart photo frame.
- Delete data:** Another person deleted data stored on one of your smart devices. This can include deleting photos, deleting saved preferences, or deleting application data.
- Change settings:** Another person changed the settings on one of your smart devices. This can include changing account settings, changing how a device behaves, changing recommendations, or logging out of an account.
- Trigger unwanted behavior:** Another person triggered unwanted behavior on one of your smart devices. This can include setting off an alarm or triggering a voice assistant at an inopportune time.
- Make purchase:** Another person used one of your smart devices to make a purchase. This can include one-time purchases or subscriptions.

In the prevalence survey, we asked participants whether they had experienced a similar scenario on one of their own devices, engaged in something similar with someone else's device, or neither. We refer to these as our two *perspectives*:

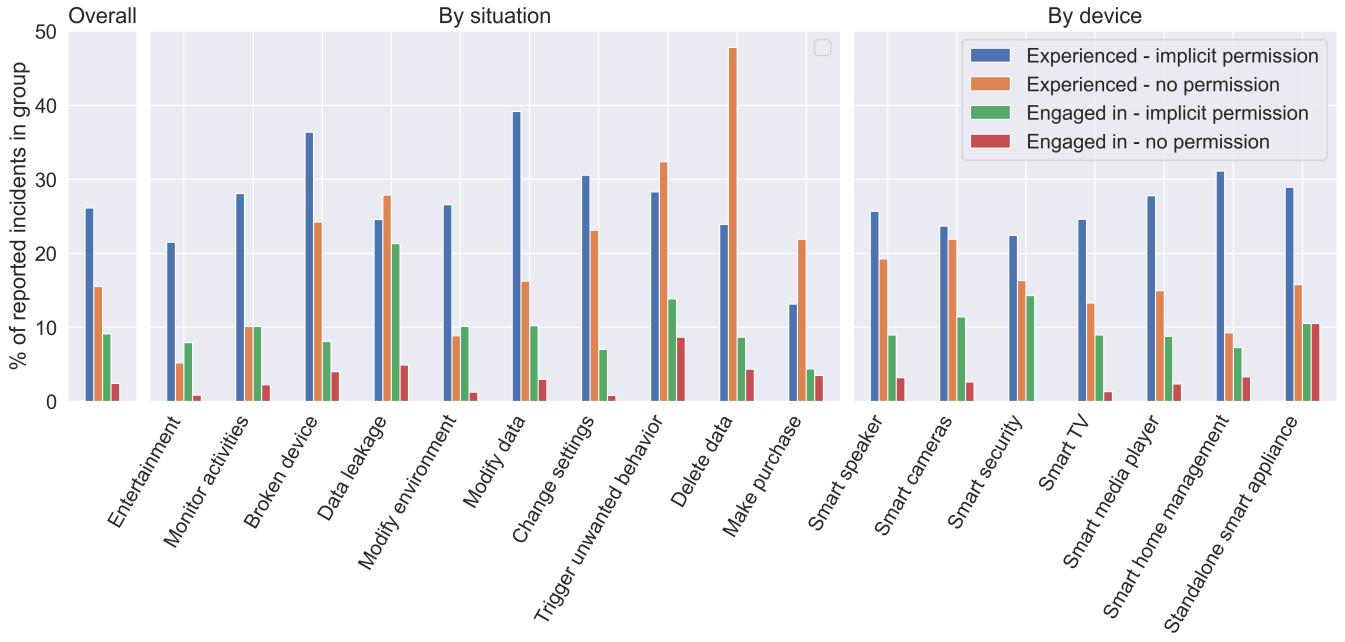


Figure 2. Proportions of unauthorized use by device and situation among all incidents reported. Proportions are calculated based on all reported incidents across all permission levels: explicit, implicit, or none (prevalence survey).

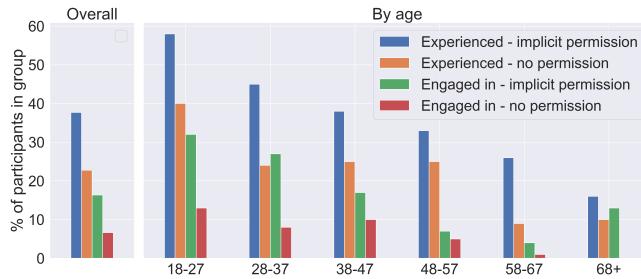


Figure 3. Proportions of unauthorized use by age groups among participants (prevalence survey).

experienced and *engaged in* respectively. For each scenario the participant identified, we then asked them to identify all devices for which this scenario applied, as well as whether they gave or received explicit permission, implicit permission, or no permission at all for that device. Figure 1 shows the number of incidents reported in the survey, organized according to the ten different scenarios, device categories, permission level (explicit, implicit, none), and perspective (experienced versus engaged in).

Of the 483 respondents to the prevalence survey, 243 reported having experienced and/or engaged in unauthorized IoT device use. Specifically, 224 participants (43%) reported experiencing unauthorized use (182 participants with implicit permission, 110 with no permission, including 68 who had experienced both). Separately, 96 (20%) participants reported engaging in unauthorized use (79 participants with implicit permission, 32 with no permission, including 15 who had

engaged in both). These numbers include 77 participants (15%) who reported both experiencing and engaging in unauthorized use.

Among unauthorized use incidents (985 incidents total), participants most commonly reported a device being used to access entertainment (246 incidents), followed by changing settings (150 incidents) and triggering unwanted behavior (145 incidents) on a device. In terms of devices, participants most commonly reported smart TVs (290 incidents), smart speakers (267 incidents), and smart media players (227 incidents) as the subject of unauthorized use. We summarize unauthorized use as a proportion of reported incidents across all permission levels, broken down by situation and by device type, in Figure 2. Implicit permission was generally more common than no permission. In contrast, for scenarios of data leakage, triggering unwanted behavior, deleting data, and making a purchase, experiences with no permission were more common than experiences with implicit permission.

Because different age groups tend to use IoT technology differently [55], we also examine unauthorized use by age group. As shown in Figure 3, the highest incidence of unauthorized use (experienced or engaged in) reported in the prevalence survey appears in the youngest age group, 18–27 years.

Many fewer participants reported engaging in unauthorized use than experiencing unauthorized use. This mismatch, observed across all participant age groups as well as all scenarios and devices, is likely due to social desirability.

In Figure 4, we look at participants who report owning at least one of a given device type and examine what proportion of them reported experiencing unauthorized use of that type

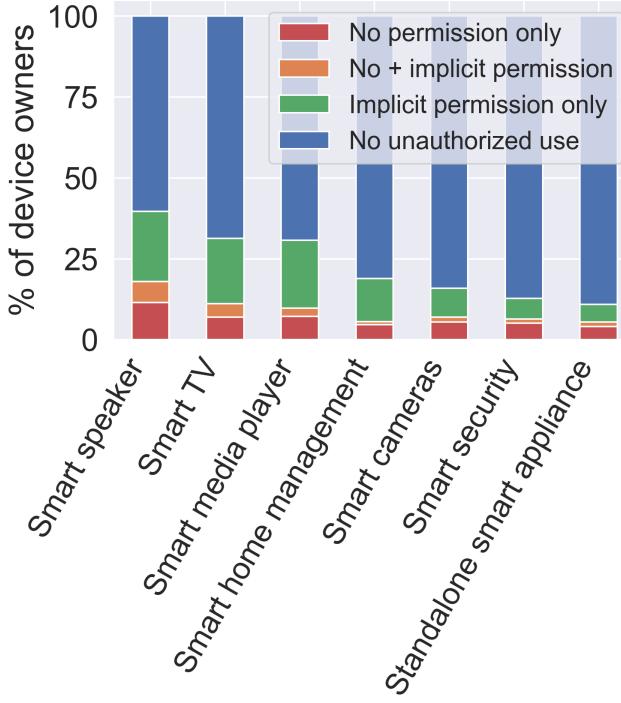


Figure 4. Proportions of participants who have experienced unauthorized use (no permission, implicit permission, or both in the cases where participants report multiple events) among participants that report owning at least one device in the given category. Device ownership counts can be found in Table 1 (prevalence survey).

of device. Among smart speaker owners, 40% report at least one unauthorized use incident, including 18% who report at least one incident involving no (as opposed to implicit) permission. Other devices with fairly high proportions of unauthorized use include smart TVs and media players. On the other hand, standalone smart appliances such as smart refrigerators and coffee makers saw the lowest rate of unauthorized use among owners (11%). Unauthorized uses of smart security devices, such as locks, were similarly uncommon.

4.2. Characterizing misuse

Of 243 in-depth responses to the prevalence survey describing specific unauthorized use incidents, 110 were classified as misuse according to the criteria in Section 3.4. Of these incidents, 91 were from the perspective of someone who experienced misuse on one of their own devices (19% of total prevalence survey participants) and 19 were from the perspective of someone who engaged in misuse on someone else’s device (4% of total prevalence survey participants). The distribution of these incidents across perspectives, devices, and scenarios is shown in Figure 5. We note that because we used dynamic weighting to select a broad range of incidents for the drill-down section, this distribution should not be

interpreted as prevalence, but rather treated as a lower bound for home IoT misuse.

We next describe qualitatively the kinds of misuse incidents participants reported in both surveys.

What types of misuse incidents happen?

Across the characterization survey and the prevalence survey, we observed misuse incidents that spanned a wide range of severity and intent. We were able to collect at least one misuse incident for every scenario category defined in Section 4.1.

Some misuse incidents were relatively benign, such as accidental disclosures of comparatively mundane information (“They accidentally said something about an Amazon cart, and it listed everything I had put in my Amazon cart the day before,” C#76) or pulling pranks (“I said, ‘Alexa, Red Alert!’ as a joke, and it started making very loud ‘red alert’ noises from Star Trek,” P#137). However, even apparently-benign events could leave device owners feeling deeply uncomfortable or that their privacy had been violated. For instance, P#355 elaborated on a time when a friend used their smart media player to access movies and changed some settings in the process, which left P#355 “angry because it felt very violating.” Similarly, P#383 reported feeling extremely uncomfortable after their smart TV had been used by a pet sitter, because they “did not think they would mess with any of my stuff.”

In addition to these relatively low-stakes incidents, we also observed less common but more severe incidents like long-term surveillance. P#151 recounted that their partner “somehow obtained the login credentials and was using the Ring doorbell to watch as I entered my home, left my home, or did any yard work outside my home. I was unaware of this surveillance until they decided to scare me by speaking through the device attached to my Ring doorbell.” This left P#151 feeling “violated and unsafe.” P#108 similarly elaborated on a time when an acquaintance used one of their devices for “listening in the room when she wasn’t present.”

While we highlight examples of intentional misuse above, we also observed incidents where misuse was perpetrated by accident. P#150 reported feeling extremely uncomfortable after “My aunt had asked [my smart speaker] for deals and triggered Alexa to purchase something” by accident when she had no permission to use the device. P#227 described a time when “A friend deleted saved media (photos and videos) stored on a Samsung Smart TV. I believe inadvertently,” leading them to feel “a little upset due to the sentimental value of the media deleted.”

Who engages in and experiences misuse? Consistent with research on social insiders and phone snooping [38, 40], the vast majority of the misuse incidents participants described involved close relations such as friends, family members, and roommates. We note the contrast between this finding and prior work on smart home security, which often emphasizes strangers, or perhaps guests or employees with short-term access [4–6]. The breakdown of relationships between participants and devices owners/users in the reported misuse cases is shown in Figure 6.

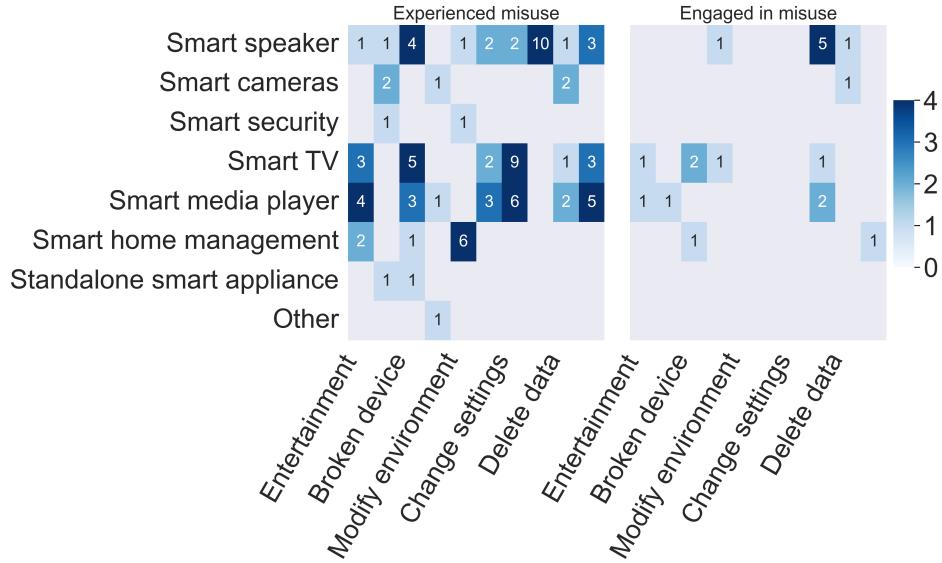


Figure 5. Counts of misuse incidents that participants elaborated on across perspectives, device types, and scenarios (prevalence survey),

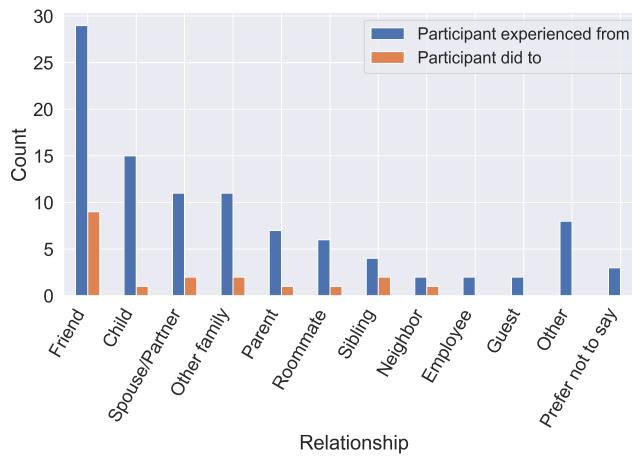


Figure 6. Reported relationships between those who experienced and engaged in misuse. Counts may not add to 110 due to multiple response (prevalence survey).

What kinds of devices are misused? We observe misuse across all the device categories listed in Figure 1.

The distribution of incidents is heavily related to both device ownership rates and our attempt to balance the scenarios and devices we asked about. Among prevalence survey participants, 74% own at least one smart TV, 63% own at least one smart media player, and 61% own at least one smart speaker. Only 16% and 15%, respectively, own smart security devices or standalone smart appliances. Accordingly, we obtained many more misuse examples for popular devices. Still, we observe that smart speakers seem particularly susceptible to triggering unwanted behavior, while changing settings is primarily a threat for smart TVs and media players.

Although we were only able to collect a few misuse

incidents for smart appliances and cameras, several of these incidents were severe: P#64 was monitored through their own smart security camera, and P#483 had their smart fridge’s camera accessed without permission.

4.3. When does unauthorized use become misuse?

Not all unauthorized use crosses the line into misuse; in the drill-down sections of both surveys, we asked detailed questions intended to distinguish whether an event qualified as misuse according to the criteria given in Section 3.4.

Though we observed many similar incidents of unauthorized use, the device owner’s comfort — and as a consequence whether an incident should be classified as misuse — varied based on personal preferences and situational context. For instance, two participants recounted incidents where unauthorized users changed the settings on the participants’ devices (a smart TV and a smart media player) to better suit their own preferences. P#325 was relatively unbothered by what happened: “It didn’t bother me much. It was explained to me that it was accidental. I believe them and no serious harm was done.” In contrast, P#473 felt that the event was an extreme violation of their boundaries: “I was upset when this occurred. I felt violated and was unable to trust the individual for what had occurred.” Similarly, while several participants recounted stories about being comfortable with housemates and guests reprogramming their smart thermostats, P#66 was displeased with one such event because “... It is my home and I choose what temperature I want the house to be since I pay the bills... I find it rude to just touch someone’s things without permission, regardless of your relationship with that person.”

Even events that appear at first glance to be severe, such as changing the owner’s personal settings or making unauthorized purchases, were considered acceptable by some

participants under certain circumstances. For example, several participants recounted incidents where their devices (and their linked payment information) were used by others to make purchases without authorization, but family ties or the amount of money spent being relatively small helped the participants feel more comfortable with these unauthorized uses. This was the case for P#168: “I was OK with it. It was an accident. It was only a \$1.99 charge.”

We qualitatively analyzed these responses and inductively identified five main factors, described below, that influenced participants’ comfort or discomfort. We initially identified these factors in the characterization survey and found that they generally held in the prevalence survey.

Owner/user relationship The relationship between the two parties was a prominent factor in whether the participant was comfortable with the unauthorized use incident they recounted. Unsurprisingly, close bonds tended to engender comfort when experiencing unauthorized use or provide justification when engaging in it. For example, P#212 said, “It didn’t bother me at all, she is like a sister to me and I know she knew how to use the device,” in regard to their smart speaker being used with implicit permission. P#141, describing a family member accessing their watch history through a smart media player, said, “It was just family, so not a big deal. Possibly more important if it’s a new friend or romantic interest.” Similarly, P#176 said their niece using their smart TV “was fine. Unexpected, but now I know that people will just use those devices without asking sometimes. So lesson learned. I would have an issue if it was not family or friends though.”

Despite never receiving explicit permission, P#31 believed that they had implicit permission to use their friend’s smart speaker, saying, “It felt natural since I always hang out with them, so I feel comfortable using their stuff.” In some cases, social bonds were used to justify using devices for pranks. C#77 explained that they were comfortable using their parents’ smart speaker for pranks because “It was in my parents home, they wouldn’t care.”

However, while closeness and trust between parties can dampen the perceived severity of an event, it can also amplify discomfort if trust is breached. For instance, P#183, whose smart media player was used without any permission, said, “I was a little disappointed with the actions of this person because, I thought we had a better relationship than that.” P#483 expressed similar sentiments after a visitor accessed their smart fridge’s camera footage without any permission, stating, “I felt like I could no longer trust this person to be alone with my things because of their ‘curiosity’ to look around and touch my stuff. It opened my eyes to how they were when they were alone and made me feel disrespected.”

Intent Lack of malicious intent was also an important factor in participants’ comfort with unauthorized uses. A friend of P#163 modified the settings on their smart TV; P#163 was “annoyed because it felt inconsiderate, but I don’t think they did it to be malicious. You’re an adult, but we also make mistakes.” In general, respondents tended to be forgiving towards events that were accidental. For instance,

P#275 said, “I was upset but not too upset because it was on accident,” describing a guest unknowingly used their smart TV to make a purchase. Participants who engaged in unauthorized use of another person’s device cited their lack of malicious intent as a justification for their own comfort with their actions. C#67 reported comfort using a host’s Google Home because they “had no ill intention of accessing something against their wishes.”

Information sensitivity Those who experienced and engaged in unauthorized access to data expressed more comfort when they did not consider the information revealed or accessed to be sensitive. However, what constitutes sensitive information varied across participants, which has been shown to be true in other contexts [56–59]. For instance, several participants in both surveys recounted events in which their watch history was accessed on a smart TV or media player, but provided differing opinions on whether they considered this information private.

User mental models We observed that device owners expressed discomfort when unauthorized use contradicted their mental models of their device’s capabilities. This often happened when someone triggered behavior the owner did not realize their device was capable of, or if the device obeyed someone it should not have, as in the case of P#200: “I was upset and uncomfortable because I was unaware things like [Alexa settings] could be changed by another person.” Other participants expressed surprise and discomfort about how easy it was for people to trigger unwanted behaviors or change settings without permission. In some cases, this weakened the owners’ trust in their devices. For instance, C#81 recalled an event making them feel “a little uncomfortable because it made me think about how smart our devices are getting.”

Severity of consequences Finally, lack of long-term consequences or harm was cited as a reason for comfort among those who experienced and engaged in unauthorized use. In particular, while many participants expressed annoyance at having to revert their devices after someone used them, they also tended to view unauthorized use events less severely if they could be easily reversed. P#260 was comfortable with their smart lights being used with implicit permission, saying, “...I could easily change my lighting back to what I prefer.” Similarly, P#346 was comfortable disabling another person’s smart home devices while house sitting in order to protect their own privacy, because “functionality was easily restored to perfect working order.”

Conversely, long-term consequences or harm resulting from unauthorized use caused discomfort among participants. P#455, whose smart speaker stopped working after their son used it, was “frustrated. I still have not worked out the reprogramming. I will do that at some time. I am not great with electronics, so it takes me forever to fix things.”

Regression on participant discomfort Using prevalence survey data, we performed an ordinal logistic regression to quantitatively surface factors that correlate with discomfort

Variable	Value	Odds Ratio	Conf. Int.	p-value
Smart device type	TV	—	—	—
	Media player	0.8	[0.3, 2.0]	0.572
	Speaker	0.7	[0.2, 1.9]	0.422
	Home management	0.9	[0.2, 4.5]	0.903
	Camera	0.3	[0.05, 1.4]	0.126
	Security	0.8	[0.04, 13.6]	0.905
	Standalone appliance	0.2	[0.02, 1.6]	0.112
Incident type	Entertainment	—	—	—
	Monitor activities	57.6	[6.9, 551.5]	<0.001*
	Broken device	22.0	[5.3, 98.9]	<0.001*
	Data leakage	1.6	[0.3, 8.8]	0.560
	Modify data	2.8	[0.7, 10.9]	0.139
	Delete data	10.1	[1.7, 64.5]	0.012*
	Modify environment	1.2	[0.3, 4.9]	0.821
	Change settings	5.1	[1.5, 17.4]	0.008*
	Trigger unwanted	1.8	[0.5, 6.3]	0.330
	Make purchase	4.1	[1.1, 16.9]	0.049*
Parent	True	2.1	[0.6, 7.6]	0.268
Child	True	0.3	[0.1, 0.8]	0.018*
Sibling	True	0.2	[0.08, 0.8]	0.017*
Spouse/Partner	True	1.5	[0.5, 4.6]	0.445
Friend	True	0.8	[0.3, 1.8]	0.561
Roommate	True	2.0	[0.6, 7.1]	0.285
Neighbor	True	0.4	[0.02, 6.01]	0.460
Employee	True	7.1	[0.5, 101.5]	0.132
Permission	Implicit permission	—	—	—
	No permission	5.5	[2.6, 11.5]	<0.001*
Frequency	One-time	—	—	—
	More than once	0.6	[0.3, 1.2]	0.132
	Not sure	21.7	[1.8, 315.6]	0.018*

TABLE 2. ORDINAL LOGISTIC REGRESSION MODEL FOR PREVALENCE SURVEY PARTICIPANTS’ DISCOMFORT WITH UNAUTHORIZED USE EVENTS. ODDS RATIOS ABOVE 1 INDICATE HIGHER DISCOMFORT RELATIVE TO THE BASELINE (“SMART TV” FOR DEVICE TYPE, “ENTERTAINMENT” FOR SITUATION, “FALSE” FOR RELATIONSHIPS, “IMPLICIT PERMISSION” FOR PERMISSION, AND “ONE-TIME” FOR FREQUENCY). PSEUDO-R²: 0.64.

among participants who had experienced unauthorized use of their smart home devices.

The dependent variable in our model was participant discomfort, and the following factors were used as potential covariates:

- **Device type**
- **Incident type**
- **Relationship of the unauthorized user to the participant** (child, parent, sibling, other family, spouse/partner, friend, roommate, neighbor, employee, or other; multi-selection possible)
- **Type of permission** (implicit, no permission)
- **Intentional** (accidental, purposeful, unsure)
- **Frequency of unauthorized use** (one-time, more than once, unsure)

We excluded one response from our regression analysis because it was the only response whose device type was “other”. We binned the “guest” relationship into the “other” relationship factor because it was only reported once.

In order to check for multicollinearity among these factors, we performed a Variance Inflation Factor (VIF) test on the initial model [60]. All variables except the “other” relationship factor scored under the VIF threshold value of 5.

Since relationship categories were each independent binary factors, rather than a single categorical choice, we excluded the “other” relationship factor from the model.

Finally, we compared a set of potential models. Every model we tested included incident type, device type, and relationship (minus the removed “other” relationship factor). We tested all possible combinations of the other covariates, but excluded interaction factors due to insufficient power. We selected our final model based on minimum Akaike Information Criterion (AIC) [50]. The final model can be found in Table 2 — this model has a Pseudo-R² of 0.64 using the Aldrich-Nelson method, as evaluated by Hagle and Mitchell [61], indicating strong fit. Odds ratios above 1 correspond to increased discomfort (more likely to be misuse) relative to the baseline.

Device type: We selected smart TVs as the baseline device type because they were the most commonly owned device in our sample. No device type exhibited significant differences in contribution to discomfort relative to smart TVs.

Incident type: We selected entertainment as our baseline scenario because we deemed it the least privacy-invasive (under typical circumstances), and it was also the most commonly reported usage scenario across all permission types (Figure 1). Among the incidents we collected, those that involved monitoring activities (odds ratio: 57.6), breaking a device (odds ratio: 22.0), deleting data (odds ratio: 10.1), changing settings (odds ratio: 5.1), or making a purchase (odds ratio: 4.1) were significantly correlated with greater discomfort. Monitoring activities and breaking a device in particular were highly associated with discomfort, with odds ratios significantly greater than other factors.

It is worth noting that while these factors make discomfort more likely, these behaviors were sometimes reported as comfortable by participants because of the factors described above, such as lack of lasting damages (P#37, whose smart TV settings were changed, said “It was no big deal, as long as it was not broke, no major issue”) or intention (P#228, whose device was used to make a purchase with their own money, was not upset because they “considered it an isolated incident. It was an accident with no foul intentions”).

Relationship: Because participants could select multiple options for relationship to cover scenarios where multiple users were involved, data types were modeled in the regression as independent boolean factors (each with baseline false). Participants reported significantly less discomfort when their device was used without authorization by their child (odds ratio: 0.3) or their sibling (odds ratio: 0.2).

Other factors: Participants reported more discomfort if they were unsure how often their device was used without authorization (compared to baseline of once; odds ratio: 21.7) and when there was no permission granted (compared to baseline of implicit permission; odds ratio: 5.5).

Despite several participants highlighting in open-response questions how intent (such as if what happened was accidental or done without malice) influenced how they felt about an unauthorized use incident, the *intent* covariate (accidental, purposeful, unsure) did not appear in the final model. This

could indicate that this factor is only important to certain participants (rather than the sample as a whole), or only in interaction with other factors, which we did not measure.

5. Discussion

Using our characterization survey and prevalence survey, we developed 10 categories of unauthorized use scenarios for smart devices, enumerated the characteristics that transform unauthorized use into misuse, and estimated a lower bound for the prevalence of misuse among a representative population sample. These unauthorized use scenarios represent everyday, interpersonal attacks that require little technical sophistication, as compared to high-effort technical attacks from remote adversaries. We discuss several key themes arising from our findings.

Local, interpersonal unauthorized use of IoT devices is common In the prevalence survey, 50% of participants had either experienced unauthorized use, engaged in unauthorized use, or both.

In addition, 19% of our survey participants report experiencing a misuse event on one of the smart home devices that they owned. This statistic should be considered a lower bound due to our study design: out of every unauthorized use incident a participant reported, they were only asked about one such event in-depth, and we used these in-depth reports to determine if misuse occurred. Thus, it is possible that some misuse events were not sampled for follow-up.

We observed misuse in every IoT device category and at least one occurrence of each unauthorized use scenario we asked about. It is important to note that essentially all the unauthorized use and misuse incidents reported by participants were low-tech, unsophisticated activities conducted primarily by family, friends and roommates, in some cases by mistake. This is consistent with previous studies on mobile phone snooping that suggest social insiders should be incorporated into adversarial models [38, 40]. With respect to IoT, researchers have thoroughly investigated important threat models like remote adversaries, data collection and misuse by manufacturers and data brokers, and sophisticated local attackers using complex technical exploits. We argue that the space of interpersonal adversaries is equally important, but less well understood, with significantly less effort put into developing solutions. In general, designing usable interpersonal access control that doesn't require implausible effort from users is a hard, unsolved problem. Previous research in the multi-user smart home setting has shown that users sometimes choose not to set explicit access controls on their devices because they trust housemates not to misuse them [8].

Nevertheless, as home IoT devices become increasingly ubiquitous, finding acceptable access-control models is a problem researchers must continue to tackle. The inherently multi-user nature of many smart home devices suggests the importance of designing interfaces that limit accidental misuse (e.g., accidental disclosure); this is also a topic ripe for further research.

In our recruitment, we did not specifically seek out adversarial households, which have different social insider threats than non-adversarial households. Smart devices may amplify harms in adversarial households, such as in the context of intimate partner violence or abuse [8]. Examining how unauthorized use and misuse of IoT devices occur in this context and their effects could inform designs to mitigate the potential harms that may arise.

Implicit permissions are important but sometimes unclear A large fraction of the unauthorized use and misuse incidents we collected involve the perception of implicit permission. In many instances, participants believed they were allowed to use a smart device because they had close ties to the device owner, even if the issue was never discussed. As P#31 described when using a smart speaker they believed that they implicitly had permission to use, “It felt natural since I always hang out with them, so I feel comfortable using their stuff.” This was even true for some participants who *experienced* unauthorized use — “I was not concerned they are good friends of ours, and we have done the same at their house,” reported P#95 in reference to a friend using their smart speaker.

However, mismatches between different people's assumptions about implicit permissions can result in misuse scenarios that create discomfort or even cause harm. For example, an owner telling a guest to “make myself at home while I was staying there,” (P#115) may create ambiguity as to whether permission to use smart devices is included. Future research could compare the perspectives of owners and non-owners to more precisely characterize conflicts in perceived implicit permission.

Misuse is highly contextual Our results are broadly consistent with, and can be understood through the lens of, the theory of contextual integrity (CI) [13]. CI sees privacy preferences as context-dependent, and that context is constituted by five factors: data subject, data type, sender, recipient, and transmission principle. We see this in our results: what the data is and who finds out are some of the most important factors considered by people. The transmission principle—most often, whether someone has permission to access a device and its data—is also crucial for people's comfort.

Not all misuse is privacy-related (e.g., breaking a device or changing its settings), but we find that contextual factors are important for adjudicating these cases as well. There has been significant prior research considering how to incorporate context into access control (e.g., [62–65]); researchers and device manufacturers should build on this work to develop access-control mechanisms that take context into account.

Misuse is highly individual We identified a variety of factors that influence (dis)comfort with unauthorized use. These factors are highly individual — superficially similar incidents can induce very different reactions. This mismatch between what individuals consider to be unacceptable may lead to scenarios of accidental device misuse, where an unauthorized user believes an action is appropriate, but the

device owner considers it unacceptable. Attempting to resolve these conflicts through general principles seems unlikely to succeed. When asked about what makes IoT use acceptable or unacceptable, many participants cited general ideas like “violation of privacy,” but in practice participants instantiated these ideas very differently, with little consistency in which actions were considered privacy violating.

This result also fits with CI, which posits that expectations are derived from social norms, which can vary in their universality and strictness, particular for nascent technologies. Some of these inconsistencies may resolve over time, as smart home devices become more common and social norms around their use become codified, but in the meantime developing technical solutions to navigate potentially conflicting expectations will be challenging. Our results provide an important reminder for device manufacturers, policymakers, and researchers to not make assumptions about acceptable practices and patterns of behaviors and instead provide support for the heterogeneity of user preferences.

As social norms stabilize, researchers should consider access-control designs that incorporate or take advantage of norms and social relationships, such as reactive access control [66] or AI systems that can detect and warn against norm-violating behaviors [8].

Misunderstanding of devices contributes to misuse

One interesting factor in whether an incident was classified as misuse was whether the incident contradicted the device owner’s preexisting mental model of the device’s capabilities. In several cases, owners expressed discomfort explicitly because unauthorized use exposed device capabilities the owner was not previously aware of. Therefore, helping users better understand their devices could help them set their expectations appropriately or even use existing tools for managing access control to better protect themselves against interpersonal misuse. This might take the form of better consumer education, clearer labels or documentation for new devices, and ideally interfaces that illustrate capabilities more intuitively. This may be challenging, especially considering the different perceptions of device owners and non-owners, all of whom may have differing levels of understanding on the potential privacy risks of IoT devices. While there is prior work on developing labels for IoT privacy through expert consultation [67], further research should evaluate this technique’s effectiveness in user education.

6. Conclusion

While home IoT devices can be targets of sophisticated external attacks, they are also easily accessed by anyone in the same physical space, such as friends, family, and housemates. They can then potentially be used in ways that cause owners discomfort or harm. We conducted two surveys to characterize the kinds of unauthorized use that people experience on their devices from others around them, understand how these incidents become misuse, and estimate the prevalence of unauthorized use and misuse.

We observe local, interpersonal unauthorized use of home IoT devices to be fairly common, having been experienced

by 43% of our prevalence survey participants. The categories of unauthorized use scenarios cover a wide range of severity, ranging from fairly innocuous events like accidental exposure of watch history to severe events like long-term spying. We show the potential for any one of these unauthorized use scenarios, even ones that seem innocuous, to become misuse incidents based on highly individual and contextual factors, such as the relationship between device owner and unauthorized user. Of our prevalence survey participants, at least 19% experienced some kind of misuse on one of their home IoT devices. Mismatches between what device owners and users believe to be acceptable behavior may be one source of unintentional misuse; differences between what permissions an owner believes they gave and what permissions a user believes they received may be another contributing factor.

Acknowledgments

We thank our participants for their contribution to our research.

This paper results from the SPLICE research program, supported by a collaborative award from the National Science Foundation (NSF) SaTC Frontiers program under award numbers 1955805 and 1955228. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of NSF. Any mention of specific companies or products does not imply any endorsement by the authors, by their employers, or by the NSF.

References

- [1] Q. Wang, P. Datta, W. Yang, S. Liu, C. Gunter, and A. Bates, “Charting the Attack Surface of Trigger-Action IoT Platforms,” in *CCS 2019*.
- [2] C. Cobb, M. Surbatovich, A. Kawakami, M. Sharif, L. Bauer, A. Das, and L. Jia, “How risky are real users’ IFTTT applets?” in *SOUPS 2020*.
- [3] S. Manandhar, K. Moran, K. Kafle, R. Tang, D. Poshyvanyk, and A. Nadkarni, “Towards a natural perspective of smart homes for practical security and safety analyses,” in *IEEE S&P 2020*.
- [4] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, “Light commands: Laser-Based audio injection attacks on Voice-Controllable systems,” in *USENIX Security 2020*.
- [5] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, “DolphinAttack: Inaudible voice commands,” in *CCS 2017*.
- [6] X. Yuan, Y. Chen, Y. Zhao, Y. Long, X. Liu, K. Chen, S. Zhang, H. Huang, X. Wang, and C. A. Gunter, “CommanderSong: A systematic approach for practical adversarial voice recognition,” in *USENIX Security 2018*.
- [7] C. Geeng and F. Roesner, “Who’s in control? Interactions in multi-user smart homes,” in *CHI 2019*.

[8] E. Zeng and F. Roesner, “Understanding and improving security and privacy in multi-user smart homes: A design exploration and In-Home user study,” in *USENIX Security 2019*.

[9] Y. Huang, B. Obada-Obieh, and K. K. Beznosov, “Amazon vs. my brother: How users of shared smart speakers perceive and cope with privacy risks,” in *CHI 2020*.

[10] M. Tabassum, T. Kosinski, and H. R. Lipford, “I don’t own the data”: End user perceptions of smart home device data practices and risks,” in *SOUPS 2019*.

[11] E. Ronen and A. Shamir, “Extended functionality attacks on IoT devices: The case of smart lights,” in *EuroS&P 2016*.

[12] J. Valente, M. A. Wynn, and A. A. Cardenas, “Stealing, spying, and abusing: Consequences of attacks on internet of things devices,” *IEEE S&P*.

[13] H. Nissenbaum, “Privacy in context,” in *Privacy in Context*. Stanford University Press, 2009.

[14] P. E. Naeini, S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L. F. Cranor, and N. Sadeh, “Privacy expectations and preferences in an IoT world,” in *SOUPS 2017*.

[15] K. Bonilla and A. Martin-Hammond, “Older adults’ perceptions of intelligent voice assistant privacy, transparency, and online privacy guidelines,” in *SOUPS 2020*.

[16] A. E. Moorthy and K.-P. L. Vu, “Privacy concerns for use of voice activated personal assistant in the public space,” *International Journal of Human-Computer Interaction*, vol. 31, no. 4, 2015.

[17] M. E. Sweeney and E. Davis, “Alexa, are you listening? An exploration of smart voice assistant use and privacy in libraries,” *Information Technology and Libraries (Online)*, vol. 39, no. 4, 12 2020.

[18] G. Germanos, D. Kavallieros, N. Kolokotronis, and N. Georgiou, “Privacy issues in voice assistant ecosystems,” in *IEEE World Congress on Services 2020*.

[19] E. Alepis and C. Patsakis, “Monkey says, monkey does: Security and privacy on voice assistants,” *IEEE Access*, vol. 5, 2017.

[20] Y. Liao, J. Vitak, P. Kumar, M. Zimmer, and K. Kritikos, “Understanding the Role of Privacy and Trust in Intelligent Personal Assistant Adoption,” in *Information in Contemporary Society*, 2019, vol. 11420.

[21] P. Cheng and U. Roedig, “Personal voice assistant security and privacy;a survey,” *IEEE*, vol. 110, no. 4, 2022.

[22] N. Malkin, J. Bernd, M. Johnson, and S. Egelman, ““What Can’t Data Be Used For?”: Privacy Expectations about Smart TVs in the U.S.” in *European Workshop on Usable Security 2018*.

[23] J. Lau, B. Zimmerman, and F. Schaub, “Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers,” *ACM Human-Computer Interaction*, vol. 2, no. CSCW, 2018.

[24] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, “Understanding the Mirai botnet,” in *USENIX Security 2017*.

[25] E. Zeng, S. Mare, and F. Roesner, “End user security and privacy concerns with smart homes,” in *SOUPS 2017*.

[26] S. Bhunia and M. Gurusamy, “Dynamic attack detection and mitigation in iot using sdn,” in *ITNAC 2017*, 2017.

[27] J. Mao, S. Zhu, X. Dai, Q. Lin, and J. Liu, “Watchdog: Detecting ultrasonic-based inaudible voice attacks to smart home systems,” *IEEE Internet of Things Journal*, vol. 7, no. 9, 2020.

[28] M. Letafati, A. Kuhestani, K.-K. Wong, and M. J. Piran, “A lightweight secure and resilient transmission scheme for the internet of things in the presence of a hostile jammer,” *IEEE Internet of Things Journal*, vol. 8, no. 6, 2021.

[29] A. A. Pammu, K.-S. Chong, W.-G. Ho, and B.-H. Gwee, “Interceptive side channel attack on AES-128 wireless communications for IoT applications,” in *APCCAS 2016*.

[30] J. Park and A. Tyagi, “Using power clues to hack IoT devices: The power side channel provides for instruction-level disassembly,” *IEEE Consumer Electronics Magazine*, vol. 6, no. 3, 2017.

[31] J. Deogirikar and A. Vidhate, “Security attacks in IoT: A survey,” in *I-SMAC 2017*.

[32] W. S. Albayaydh and I. Flechais, “Exploring bystanders’ privacy concerns with smart homes in jordan,” in *CHI 2022*.

[33] C. Cobb, S. Bhagavatula, K. A. Garrett, A. Hoffman, V. Rao, and L. Bauer, ““I would have to evaluate their objections”: Privacy tensions between smart home device owners and incidental users,” *PETS 2021*.

[34] K. Marky, N. Gerber, M. G. Pelzer, M. Khamis, and M. Mühlhäuser, ““you offer privacy like you offer tea”: Investigating mechanisms for improving guest privacy in IoT-equipped households,” *PETS 2021*.

[35] Y. Yao, J. R. Basdeo, O. R. McDonough, and Y. Wang, “Privacy perceptions and designs of bystanders in smart homes,” *ACM Human-Computer Interaction*, vol. 3, no. CSCW, 2019.

[36] S. Mare, F. Roesner, and T. Kohno, “Smart devices in Airbnbs: Considering privacy and security for both guests and hosts,” *PETS 2020*.

[37] J. Bernd, R. Abu-Salma, and A. Frik, “Bystanders’ privacy: The perspectives of nannies on smart home surveillance,” in *FOCI 2020*.

[38] D. Marques, I. Muslukhov, T. Guerreiro, L. Carriço, and K. Beznosov, “Snooping on mobile phones: Prevalence and trends,” in *SOUPS 2016*.

[39] D. Marques, T. Guerreiro, and L. Carriço, “Measuring snooping behavior with surveys: It’s how you ask it,” in *CHI EA 2014*.

[40] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov, “Know your enemy: The risk of unauthorized access in smartphones by insiders,” in *MobileHCI*

2013.

[41] A. O. Arikewuyo, K. K. Eluwole, and B. Özad, "Influence of lack of trust on romantic relationship problems: The mediating role of partner cell phone snooping," *Psychological Reports*, vol. 124, no. 1, 2021.

[42] A. K. Karlson, A. B. Brush, and S. Schechter, "Can I borrow your phone? Understanding concerns when sharing mobile phones," in *CHI 2009*.

[43] J. R. Frampton and J. Fox, "Monitoring, creeping, or surveillance? A synthesis of online social information seeking concepts," *Review of Communication Research*, vol. 9, 2021.

[44] S. T. Hawk, A. Becht, and S. Branje, "“Snooping” as a distinct parental monitoring strategy: Comparisons with overt solicitation and control," *Journal of Research on Adolescence*, vol. 26, no. 3, 2016.

[45] H. Carter, "Misuse of library public access computers," *Journal of Library Administration*, vol. 36, no. 4, 2002.

[46] V. Oksman, *The mobile phone – A medium in itself*. VTT, 2010.

[47] J. Tang, E. Birrell, and A. Lerner, "Replication: How well do my results generalize now? The external validity of online privacy and security surveys," in *SOUPS 2022*.

[48] J. Saldaña, *The coding manual for qualitative researchers*, 2021.

[49] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, no. 1, 1977.

[50] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Selected papers of hirotugu akaike*. Springer, 1998.

[51] P. Grimm, "Social desirability bias," *Wiley international encyclopedia of marketing*, 2010.

[52] E. M. Redmiles, S. Kross, and M. L. Mazurek, "How well do my results generalize? Comparing security and privacy survey results from MTurk, web, and telephone samples," in *IEEE S&P 2019*.

[53] E. Peer, D. Rothschild, A. Gordon, Z. Evernden, and E. Damer, "Data quality of platforms and panels for online behavioral research," *Behavior Research Methods*, 2021.

[54] P. I. Fusch and L. R. Ness, "Are we there yet? Data saturation in qualitative research," *The qualitative report*, vol. 20, no. 9, 2015.

[55] A. J. A. M. van Deursen, A. van der Zeeuw, P. de Boer, G. Jansen, and T. van Rompay, "Digital inequalities in

[58] R. Wash and E. Rader, "Too much knowledge? Security beliefs and protective behaviors among united states the internet of things: differences in attitudes, material access, skills, and usage," vol. 24, no. 2, 2021.

[56] N. Warford, C. W. Munyendo, A. Mediratta, A. J. Aviv, and M. L. Mazurek, "Strategies and perceived risks of sending sensitive documents," in *USENIX Security 2021*.

[57] S. Ruoti, T. Monson, J. Wu, D. Zappala, and K. Seamons, "Weighing context and trade-offs: How suburban adults selected their online security posture," in *SOUPS 2017*. internet users," in *SOUPS 2015*.

[59] A. Woodruff, V. Pihur, S. Consolvo, L. Brandimarte, and A. Acquisti, "Would a privacy fundamentalist sell their DNA for \$1000... if nothing bad happened as a result? The Westin categories, behavioral intentions, and consequences," in *SOUPS 2014*.

[60] T. A. Craney and J. G. Surles, "Model-dependent variance inflation factor cutoff values," *Quality engineering*, vol. 14, no. 3, 2002.

[61] T. M. Hagle and G. E. Mitchell, "Goodness-of-fit measures for probit and logit," *American Journal of Political Science*, 1992.

[62] P. Klemperer, Y. Liang, M. Mazurek, M. Sleeper, B. Ur, L. Bauer, L. F. Cranor, N. Gupta, and M. Reiter, "Tag, you can see it! Using tags for access control in photo sharing," in *SIGCHI conference on human factors in computing systems 2012*.

[63] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov, "Dynamically regulating mobile application permissions," *IEEE S&P 2018*.

[64] Y. J. Jia, Q. A. Chen, S. Wang, A. Rahmati, E. Fernandes, Z. M. Mao, and A. Prakash, "ContexIoT: Towards Providing Contextual Integrity to Appified IoT Platforms," in *Network and Distributed System Security Symposium 2017*.

[65] H. R. Lipford, G. Hull, C. Latulipe, A. Besmer, and J. Watson, "Visible flows: Contextual integrity and the design of privacy mechanisms on social network sites 2009," in *2009 International Conference on Computational Science and Engineering*.

[66] M. L. Mazurek, P. F. Klemperer, R. Shay, H. Takabi, L. Bauer, and L. F. Cranor, "Exploring reactive access control," in *CHI 2011*.

[67] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, "Ask the experts: What should be on an IoT privacy and security label?" in *IEEE S&P 2020*.