End User Security & Privacy Concerns with Smart Homes

Eric Zeng, Shrirang Mare, Franziska Roesner
Paul G. Allen School of Computer Science & Engineering, University of Washington
{ericzeng,shri,franzi}@cs.washington.edu

ABSTRACT

The Internet of Things is becoming increasingly widespread in home environments. Consumers are transforming their homes into smart homes, with internet-connected sensors, lights, appliances, and locks, controlled by voice or other user-defined automations. Security experts have identified concerns with IoT and smart homes, including privacy risks as well as vulnerable and unreliable devices. These concerns are supported by recent high profile attacks, such as the Mirai DDoS attacks. However, little work has studied the security and privacy concerns of end users who actually set up and interact with today's smart homes. To bridge this gap, we conduct semi-structured interviews with fifteen people living in smart homes (twelve smart home administrators and three other residents) to learn about how they use their smart homes, and to understand their security and privacy related attitudes, expectations, and actions. Among other findings, we identify gaps in threat models arising from limited technical understanding of smart homes, awareness of some security issues but limited concern, ad hoc mitigation strategies, and a mismatch between the concerns and power of the smart home administrator and other people in the home. From these and other findings, we distill recommendations for smart home technology designers and future research.

1. INTRODUCTION

Anticipated by researchers for some time now, the Internet of Things (IoT) has arrived in the homes of end users. By some estimates, there are already hundreds of millions of connected "smart home" devices in more than 40 million homes in the U.S. alone, and by 2021, that number is expected to double [48, 56]. With the rise of consumer smart home platforms like Samsung SmartThings [55], Apple Homekit [5], and others, as well as connected devices like Amazon Echo [4], Google Home [29], and Philips Hue lightbulbs [45], end users are empowered to set up their own connected, automated, smart homes. These smart homes support desirable features, such as voice-controlled lights and remote-controlled door locks, but they also raise new security and privacy risks.

Indeed, computer security researchers have already identified numerous issues with smart home technology. These issues range from over-privileged applications running on smart home platforms [26] to viral attacks that can spread between infected lightbulbs [50]. The recent Mirai malware — which compromised connected devices and conscripted them into a botnet, disrupting the internet for millions of people [43] — shows that these risks are already leading to concrete attacks. We discuss additional examples in Section 2.

However, despite an increased focus on smart home security, and the reality of the emerging risks, there has been little study of the security and privacy concerns of *end users* who set up and use these smart home platforms and devices. Without an understanding of the concerns, needs, and use cases of these end users, researchers and smart home platform designers can neither prioritize which problems to focus on, nor develop effective solutions.

We aim to bridge this gap in this work, asking questions such as: how and why do people use their smart homes? What sorts of mental models have users developed for smart homes? What are their security and privacy concerns (or lack thereof), and how do these compare to the risks identified by security researchers? What sorts of issues play out in homes with more than one user? What security or privacy mitigation strategies do end users already use, and where are additional technical solutions or other design efforts needed?

We explore these questions through in-depth, semi-structured interviews with fifteen participants. All participants live in smart homes: twelve administer their smart homes, and three live in a smart home administered by someone else.

We find that our interview participants have an assortment of (generally sparse) threat models, and that the sophistication of their threat models often depends on their technical knowledge of smart homes. And while participants identified security and privacy issues such as data collection, surveillance, or hacking, most were not concerned about these issues on a day-to-day basis. We also identify tensions that can arise in smart homes with multiple users, which in the extreme could lead to potentially dangerous situations if the administrator of the smart home uses the technology to spy on or deny access to other users.

From our findings, we distill lessons and recommendations for future smart home platforms and devices. For example, we recommend further studying and designing consciously for multi-user interactions in smart homes, and we recommend improving user awareness and control through careful UI/UX design, including the inclusion of physical controls on devices. Ultimately, better understanding end users will help us identify gaps between current system designs and users' security needs and expectations, as well as tensions between users' functionality and security needs, and will help focus the efforts of the research community and industry.

In summary, our contributions include:

1. We conduct in-depth, semi-structured interviews with fifteen smart home users, studying how and why they use smart home technologies, their mental models, their security and privacy concerns (or lack thereof), and the mitigation strategies they employ.
2. Among our findings, we learn that participants' threat models are sparse and depend on the sophistication of their technical mental models, that many current smart

Symposium on Usable Privacy and Security (SOUPS) 2017, July 12–14, 2017, Santa Clara, California.

home users are aware of potential security and privacy issues but not generally concerned, and that tensions may arise between multiple residents in a smart home.

3. From these findings, we distill recommendations for the designs of future smart home platforms and devices, as well as identify opportunities for future work.

2. BACKGROUND AND RELATED WORK

The Internet of Things (IoT) is a broad term for internet-connected devices, which has come to encompass everything from connected cars, wearables, and connected industrial/manufacturing equipment. Our focus is on *smart home* technology, which we consider to include internet-connected appliances, lighting, sensors, door locks, and other objects designed for the home environment. This technology enables applications like security systems and remote monitoring, lighting and climate control that adapts to a user's presence and habits, and voice controls for lighting and appliances.

Current Smart Home Technology Landscape. In recent years, we have seen a rapid increase in the number and type of consumer-oriented internet-connected devices for automating home environments. While home automation technology has existed for decades, smart home devices are now internet connected, interoperable between different vendors, and controllable via smartphone.

Standalone smart devices include thermostats (e.g., Nest), lights (e.g., Philips Hue), motion detectors, door/window sensors, air quality sensors, power outlets, and door locks. Some of these devices connect to the internet through existing Wi-Fi networks, while others use low energy protocols like Zigbee and Z-Wave, and communicate to the internet through a bridge. Smart devices allow users to *automate* their home, e.g., automatically adjusting the thermostat, or turning on or off lights based on motion sensor readings.

Two types of smart home platforms have emerged: *hubs* and *cloud-based integrations*. Hubs — such Samsung SmartThings [55], Wink [2], and Vera [1] — are central hardware devices that other smart home devices communicate with, and can act as a Z-Wave or Zigbee bridge. Via the hub's companion app or website, users can program *automations*. Some hubs, like Samsung SmartThings, support *third-party apps*, which are prepackaged, complex automations written by other developers. Similar to hubs, emerging intelligent personal assistants, like the Google Home and Amazon Echo, can be integrated with many existing smart home devices, allowing users to control their smart home using their voice.

On the other hand, cloud-based integrations rely on the fact that for many stand-alone devices, commands from a user's phone to the device transits the cloud. These cloud services often expose APIs for controlling devices over HTTP. Middleman cloud services like IFTTT (If This Then That) and Stringify can use these APIs to connect stand-alone devices together, and to run automations.

Smart Home Security and Privacy Concerns. Security experts have raised concerns about the security and privacy risks with internet-connected devices in homes [6, 30, 53]. Concerns include privacy risks due to pairing and discovery protocols that leak information about devices in the home [62], insecure communication leaking sensitive information about the home and the residents [17], and vulnerabilities in the devices that can allow an attacker to remotely spy on residents or disrupt their lives [21, 22, 44]. Technological solutions when not implemented correctly may amplify social issues [58]. Shared in-home devices presents new access control challenges [59], which, if not addressed carefully, may amplify interpersonal issues among residents.

Researchers have begun analyzing smart home platforms and devices (e.g., [24, 26, 44]). Findings include over-privileged applications on smart home platforms and vulnerable devices like locks [32] and lightbulbs [42, 50]. Attacks have also occurred in the wild: the massive Mirai DDoS botnet attack disrupted the internet for millions of users [43], a glitch in the Nest thermostat left users in the cold [8], a baby monitor was hacked and a vulnerability in Foscam cameras left thousands of users vulnerable to similar attack [31], and recent reports suggest that internet-connected smart TVs can be used to record conversations [52]. Furthermore, a recent report indicates that IoT malware and ransomware attacks are on the rise [38]. In response to these concerns, researchers have begun to develop designs for more secure smart home platforms (e.g., [27, 54, 63]).

End-User Studies. Prior research on end users of smart homes has generally not focused on security and privacy issues but rather on usability issues, such as installation, motivations and use cases, and the interfaces for control and automation. [10, 20] Research in this area has identified tensions that arise due to differences between members of the household. Brush et al. and Mennicken et al. found that there is often one user who is most enthusiastic and others who interact with the smart home more passively [9, 41]. Ur et al. studied differences in privacy attitudes between teens and parents regarding home-entryway surveillance [60]. Mennicken et al. implemented a calendar based interface for smart home configuration to make it more accessible to passive users [40]. Our work surfaces a similar dynamic between primary and incidental smart home users.

Some prior work has also investigated security and privacy concerns of end users. Brush et al. [9] visited 14 smart homes to study adoption issues, and among their findings, found concerns about security-critical devices like smart door locks and cameras. Worthy et al. [61] asked five subjects to keep an ambiguous IoT device in their homes for a week, finding trust as a critical factor in IoT technology acceptance. Choe et al. [12] asked 22 participants to take devices home for four weeks and studied their perceived benefits and concerns, finding more concern than we do in our study.

Our research contrasts with prior work in three ways: first, we interview participants who have been living in a smart home for months, past the novelty phase and into day-to-day use. Second, we focus primarily on security and privacy, rather than general usability issues. And lastly, we contribute an updated understanding of usability, security, and privacy issues for the current generation of smart home devices, such as Samsung SmartThings, Amazon Echo, and Philips Hue.

Further afield, others have studied security and privacy concerns of end users for related technologies, including parent-child interactions with connected toys [39], security and privacy issues with household robots [11, 22], access control challenges in the home [37], and privacy issues with using smart home technology for assisting senior citizens [16, 57].

3. RESEARCH QUESTIONS

To inform the design of more secure smart homes in the future, we set out to investigate the following research questions.

General Smart Home Use. We ask: What are the common use cases for smart homes today? While the types of home IoT devices have proliferated in recent years, ranging from smart egg trays to smart dolls, learning which types of devices, platforms, and automations are typically present in smart homes will help us understand which security and privacy issues are most salient in this space, and which functionality or other factors are critical to users.

Smart Home Technology Mental Models. We ask: What mental models do users have of their smart home? For example, do their mental models include communication between devices in the home, and/or communication beyond the home (i.e., in the cloud)? Prior work has found that incomplete mental models about a technology leads to incomplete threat models and limited adoption or use of security tools (e.g., email encryption [49], internet privacy [34]).

Smart Home Threat Models. We aim to learn about the specific threat models and security concerns — or lack thereof — of smart home end users. Experts have developed extensive threat models for IoT and smart homes, informed by a technical understanding of the potential vulnerabilities. End users may develop different threat models. We investigate the potential gap between a security expert's threat model and what users are concerned about. What risks are users unaware of or unconcerned about, and are experts considering all of the issues that matter to end users?

Mitigation Strategies. As part of studying end user threat models, we also investigate any mitigation strategies they use when they do have security or privacy concerns. For example, do users change their in-home behaviors around devices that record audio or video? If they employ technical mitigation strategies, are these strategies actually effective?

Multi-User Interactions. What unique security or privacy issues arise in smart homes due to their shared nature? Today, people increasingly use personal computing devices that are not shared with others, like laptops or smartphones [35]. However, smart home technologies are located in common spaces and are critical to basic functions of the home, such as lighting or physical access, thereby affecting all residents. We explore whether incidental users of smart homes, who were not primarily involved in the system's configuration, hold different security and privacy concerns than the primary user, or view the primary user as a potential adversary.

Other Constraints and Requirements. In addition to security and privacy factors, we anticipate that participants will make choices about whether and how to set up their smart homes based on other factors, including convenience, functionality, usability, reliability, and latency. These constraints and requirements may affect what security and privacy solutions are acceptable for end users.

Recommendations for Researchers and Smart Home Designers. Through this investigation, we aim to develop recommendations for smart home designers and for researchers. Specifically: Where should the computer security community focus its efforts? Given the range of potential issues to address, what type of work should be prioritized, and how? For example, should we prioritize better protecting users from malicious or misbehaving third party automations? How can we design devices to promote better mental models and security behaviors? We return to these questions in Section 6.

4. METHODOLOGY

In this section, we describe our study methods and materials.

4.1 Pilot Interviews

Before designing our interview questions, we conducted an exploratory interview with a colleague who set up and lives in a smart home. After designing the initial interview questions, we conducted four additional pilot interviews with smart home residents, and made modifications to the questions to improve their clarity, and to better answer our research questions. We do not include exploratory or pilot interview data in our general results, though we present one particularly relevant anecdote from one of these interviews.

4.2 Recruitment and Screening

We recruited participants by advertising on relevant mailing lists, on smart home related Reddit communities, and via the researchers' social media accounts (Twitter, Facebook).

Potential participants were asked to fill out a screening survey, selecting which, if any, smart home platforms or devices they own, how long they have been using their smart home, whether they set it up themselves, how many other people live in the home, as well as demographic information (age, gender, profession). Participants were also asked to provide their name and email address if they were willing to participate in a phone or Skype interview. We used the screening responses to select participants with at least one smart home platform and covering a range of technical skill levels (inferred from profession); we also explicitly recruited and selected participants who used but did not set up or manage their own smart home.

Participants who completed the phone or Skype interview were compensated with a $10 Amazon gift card; participants who filled out only the survey did not receive compensation.

4.3 Interview Procedure

Participants who were selected for the full interview were then contacted by the researchers to schedule a phone or Skype call. Interviews were conducted by two researchers: one leading the interview and another taking notes and recording the session. We asked participants about:

General Questions: We asked participants to describe the smart home devices they own, how they use them, what apps or automations they have installed, and whether they access these devices remotely or only while physically in the home.

Mental Models: To elicit participants' mental models and degree of technical understanding of their smart home, we asked them to explain how their smart home works, verbally and through a drawing exercise. Drawings have been found to be an effective method for externalizing mental models in conjunction with verbal reports [33], and has been used in several studies of the relationship between mental models and security [34, 47, 49].

We allowed participants to either create a diagram electronically using Google Drawings, or to draw on paper and send us a photograph. We show examples in Section 5.

Security Concerns: In order to avoid prompting participants to merely agree with the interviewer that security and privacy concerns might arise with smart homes (i.e., avoid participant response bias [7, 19]), we began by asking more general questions that could elicit security or privacy concerns but did not explicitly mention them. We asked whether they had hesitations about getting any of their smart home devices, whether there were any devices they thought about getting but ultimately decided against, or whether there were any devices they used but later deactivated.

For participants who did not organically bring up security or privacy concerns, we then prompted specifically about security and privacy (making it clear that a lack of such concern was a valid response, again to avoid participant response bias). We also asked if they had heard about security and privacy concerns with smart homes in the news, and whether they shared those concerns or felt they were overblown; and we asked participants to compare their concern about smart homes to their concern about phones and laptops.

Mitigation Strategies: We asked participants whether their security and privacy concerns (if any) had caused any changes in behavior (e.g., acting differently around smart home devices or changing device settings).

Multi-User Scenarios: We asked participants how many people live in their home, who has what types of access to the smart home, whether they have had disagreements with others about the smart home, and whether house guests have interacted with the smart home.

Technical Skill: We asked participants to self-report, on a scale of 1 (novice) to 5 (expert), familiarity with technology in general, smart home technology, and computer security.

Wrap-Up: Finally, we asked participants if there were any questions they expected us to ask, and gave them a chance to tell us anything else about their smart home.

As an in-depth, qualitative interview, we tailored our questions to the context of individual conversations. Thus, although all participants were asked the above questions, we also asked relevant follow-up questions where appropriate. A copy of the interview protocol is provided in Appendix A.

4.4 Data Analysis

We used a bottom up qualitative method to analyze the data. Three researchers independently read notes from the interviews and listened to recordings, and generated list of themes. Then, the researchers met in person to consolidate the most salient themes into a shared codebook, which consisted 16 structural codes (based on our research questions), further divided into 116 subcodes. The structural codes were broad categories, such as "Mitigation Strategies", and the subcodes enumerated specific instances mentioned by participants, e.g. "Network segmentation". Then, each interview was independently coded by two of the three researchers. One researcher was the primary coder, and participated in coding each interview. After all interviews were coded, the researchers resolved disagreements resulting from human error or misunderstanding of the codes, where possible. Cohen's kappa, a measure of inter-coder agreement, was 0.96. (Fleiss rates kappa values over 0.75 as excellent agreement [28].) Since there are some remaining disagreements, in Section 5, we report numerical values based on the primary coder.

4.5 Ethics

This study was reviewed by our institution's IRB, and was considered exempt. We did not ask participants to reveal sensitive information like account names or home addresses. All participants provided informed consent to participate in the study and be audio-recorded. We stored all interview recordings in password-protected form and removed any identifying information from notes and transcripts.

5. RESULTS

We now turn to a discussion of our results, organized according to the research questions presented in Section 3.

5.1 Participants

Thirty-three participants completed the pre-screening survey, and we conducted interviews with 15 of them, selecting people with smart home platforms and devices, and covering a range of technical skills and other factors. Interviews were conducted in Feb. 2017 and lasted on average 38 minutes.

Of the 15 participants (summarized in Table 1), four were women, eight did not mention having a background in IT or computer science, and two were aged 55 years or older. Participants had smart homes for at least two weeks and up to eight years. Table 1 presents self-reported familiarities with technology, security, and smart homes. However, in some cases, these self-estimates seemed miscalibrated. For example, one participant reported only a "3" in technology familiarity, but was able to describe a cloud-based client-server architecture for smart homes, while others who reported high familiarity with security did not articulate specific concerns even when directly asked. Nevertheless, we include these values as rough indicators of confidence in their abilities.

5.2 General Smart Home Use

We begin by describing the smart home devices participants own, how they use these devices, and how they orchestrate automations between the devices. These details will provide context for subsequent results, and they highlight use cases that computer security solutions must take into account.

Devices. Participants reported having a large variety of internet-connected devices, from many different manufacturers. We summarize these devices in Table 2. Most common are smart lights, thermostats, cameras, and switches. Participants using their smart homes as security systems typically had sensors on doors and windows, as well as motion sensors.

Nine participants mentioned having a hardware hub, and a few others mentioned using apps for centralized control, like Apple Homekit. Intelligent personal assistants, such as the Amazon Echo or Google Home, are also very common (13).

Some of the more uncommon smart devices were custom-made by the participants. For example, P8 was able to automate the lights and jets on their swimming pool by integrating its control systems with a Raspberry Pi, and implemented custom software to decode the data stream and integrate it with their SmartThings hub.

Use Cases. We identified four common smart home use cases: increasing physical safety (including security systems, door locks, and smoke detectors; 9 participants), home automation (automatically adjusting lighting, temperature, or other devices; 13 participants), remote control, and in-home sensing. Many participants mentioned multiple use cases.

ID	Gender	Age	Profession	Primary User?	CS/IT Background?	Self-Reported Familiarity with... Technology	Computer Security	Smart Homes
P1	Male	35-44	IT Security	Yes	Yes	5	5	3-4
P2	Male	35-44	Marketing	Yes	Yes	4	4	3
P3	Female	55+	Biologist	Yes	No	2-5	3.5	3.5
P4	Male	25-34	Healthcare IT	Yes	Yes	4	4	3
P5	Male	25-34	IT Technician	Yes	Yes	5	4	4
P6	Male	25-34	Engineering PM	Yes	Yes	5	4	5
P7	Male	25-34	Fundraiser in higher ed	Yes	No	4	3	5
P8	Male	45-54	Software Engineer	Yes	Yes	5	4	4
P9	Male	25-34	Finance	Yes	No	4	4	3
P10	Male	55+	Chief Financial Officer	Yes	No	4	4	4-5
P11	Male	55+	Professor	Yes	No	3	3	3
P12	Male	18-24	Retail supervisor	Yes	No	5	4	4
P13	Female	18-24	Student	No	Yes	3	2	2-3
P14	Female	25-34	Academic Admin	No	No	2	1	2
P15	Female	18-24	Student	No	No	3-4	3	3

Table 1: Summary of participants. Familiarity was self-reported on a scale of 1 (low) to 5 (high).

Type of device	Count	Examples
Lights	15	Philips Hue, Belkin Wemo Link, Osram Lightify, HomeBrite, LIFX
Intelligent Personal Assistant	13	Amazon Echo, Google Home
Thermostat	12	Nest Thermostat, Emerson Sensi, Ecobee Thermostat
Camera	11	Nest Cam, Withings Home, Foscam, Ubiquiti Aircam
Power outlets and switches	10	Belkin Wemo, Lutron Caseta
Motion Sensor	10	—
Hub	9	Samsung SmartThings, openHAB, Vera, Abode
Door Lock	7	Kwikset Smart Lock
Smoke detector	4	Nest Protect
Leak detector	2	—

Table 2: Devices owned by participants. Only devices owned by more than one participant are listed.

Though remote device use opens the door for security and privacy risks, we find that it is a critical feature for many users: nine participants remotely controlled devices, like lights or thermostats, while eleven used devices to remotely sense within the home, including monitoring things like camera feeds, air quality, and status of devices.

Modality. Participants interact with smart home devices in several different ways, often in combination. 14 of 15 participants use a smartphone app to control or program their devices. 13 participants use contextual triggers, i.e., behavior that executes based on the context, like the time of day or whether the user is home. 12 participants use an Amazon Echo or Google Home to control their devices via voice. 8 participants mentioned using motion sensors. Some remarked that using mobile apps was tedious, and preferred to use voice controls or automations exclusively.

Automations. We define *automations* to be programs that cause devices to do something on their own, or programs that connect two different types of devices so that one can trigger the other (e.g., enabling voice-controlled lights by integrating them with the Amazon Echo). Furthermore, we distinguish between three types of automations: end user programming, custom scripting, and third party apps.

Most standalone devices and hubs feature an end user programming interface, which allows users to program automations for their home on a graphical interface, usually in a mobile app. For example, the SmartThings mobile app allows users to program "routines" for devices like lights: users can trigger lights to turn on and off based on activity from motion sensors, door sensors, time of day, or whether their phone is present in the house. We found end user programming to be the most common method for automations; 11 of 15 participants used this type of interface.

Four more technically skilled users automated their homes by writing scripts for Raspberry-Pi based controllers, like openHAB or HomeAssistant. Three others used custom scripts written by *others*: P7 and P10 downloaded scripts from smart home forums, and P14's openHAB was programmed by her husband. P10 was actually able to request others to write Vera automations for him, and when we asked him about it (incredulously), he said, "Yeah, isn't that great? I've done it 3 or 4 times." Though code taken directly from others may pose security risks, he was not concerned about this risk, as we discuss further below.

Devices can also be automated by third-party tools, such as apps on appified platforms like SmartThings, or cloud-based tools like IFTTT. These methods are used (1) to provide complex automations not possible through end user programming, like adjusting the thermostat based on outdoor temperature, and (2) to integrate devices that are not built-in to a platform, like connecting an Echo to SmartThings.

We found that third-party automations were less common than custom programming solutions. Four participants mentioned cloud services like IFTTT and five mentioned using app-based automations. Both were mostly used when hubs did not provide sufficient integration or functionality with certain device families. Two non-integration automations mentioned were a disco light app for Philips Hue, and a door lock code management app for SmartThings. As we discuss further in Section 6, this finding suggests that research efforts focusing on the security of smart home applications (e.g., [26]) may be considering only a narrow use case.

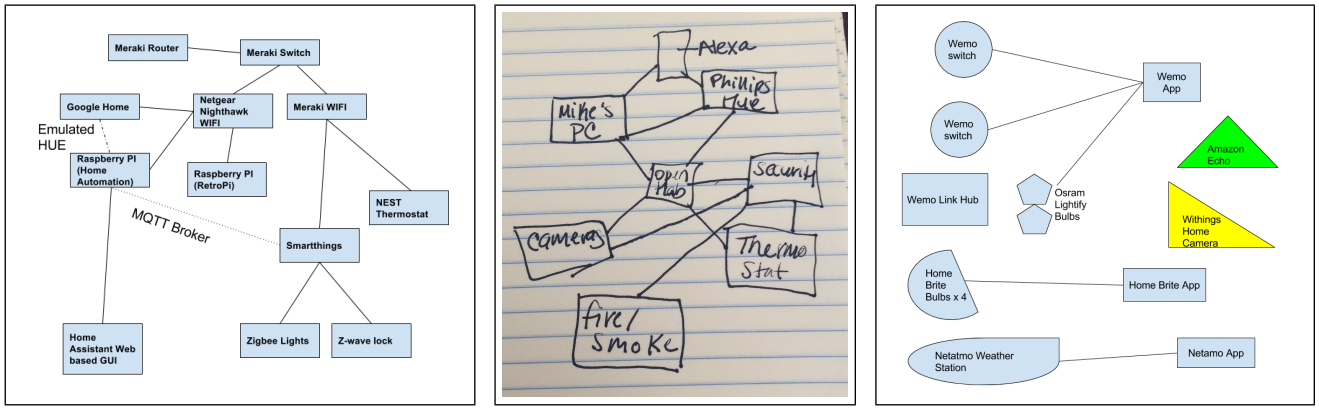

Figure 1: Participant drawings showing examples of (a) advanced (from participant P2), (b) intermediate (from participant P14), and (c) limited technical mental models (from participant P3). P2's diagram (a) shows how they used network segmentation to separate their smart home devices from their other computers. P14's diagram (b) does not represent the network topology, but rather links in functionality. In P3's diagram (c), lines are drawn between devices and their associated apps, but no technical details are captured. This diagram was edited for clarity, removing only text describing the functions of the devices.

5.3 Smart Home Technology Mental Models

Based in part on prior work linking limited technical mental models with limited adoption of security tools and incomplete threat models [34, 49], we sought to understand participants' general mental models about their smart homes before diving into security specific questions. We categorized the sophistication of participants' mental models based on both their drawings and their verbal explanation of their smart home system. Our analysis was based on codes for whether the participant demonstrated an understanding of specific technical elements of their smart home, which we describe below.

Participants with the most advanced mental models had a highly technical level of understanding of their smart home system, and were able to represent the network topology, including wireless protocols, hubs, routers, and sometimes the role of cloud servers. One example of this is P2, who was able to produce an accurate network diagram (see Figure 1a), and raised concerns about how commands traveling to the cloud affect latency. Participants in this category generally had a background in IT or computer science.

Participants with an intermediate level mental model had some sense of which devices in their home communicate with each other, but without a deep understanding of how. These users were typically capable users of technology, but did not have technical training. One participant in this category (P14) diagrammed functional relationships between devices in her home (see Figure 1b), such as between the Amazon Echo and Philips Hue lights, but did not capture the role of the cloud or their wireless router.

The last category encompasses participants who had a limited understanding of smart home technology in general, and indicated no awareness of technical details, like their network or the cloud. When we prompted to draw a diagram of his smart home system, P11 drew the physical layout of his home, and the locations of the devices, but did not illustrate how the lights and the Echo communicated with each other. Another example was P3 — in her diagram, each device had a line drawn to a shape representing the smartphone app associated with the device (see Figure 1c).

Asset	Concerned	Mentioned but not concerned
Physical security	11/15	1/15
Audio logs	4/15	4/15
General home privacy	5/15	1/15
Behavior/presence logs	2/15	2/15
Personally identifiable info	2/15	1/15
Bandwidth	1/15	0/15
Money	1/15	0/15
No identified assets	1/15	

Table 3: Assets identified by participants.

As we will see in the next section, the sophistication of a participant's technical mental model often affects the sophistication of the resulting threat model.

5.4 Smart Home Threat Models

We now turn to a core component of our study: participant threat models and security/privacy concerns (or lack thereof).

Overall, we found that participant threat models were sparse. Participants mentioned a diverse set of potential security and privacy issues, but few concrete concerns were articulated by a majority of participants. Moreover, participants were sometimes aware of potential issues but were explicitly not concerned about them. Thus, we coded threat model themes as "mentioned", "not mentioned", and "mentioned but not concerned". We summarize participant threat models in Tables 3-6, organized into assets, adversaries, vulnerabilities, and threats that came up during the interviews.

Assets. The most common asset identified by participants was physical security. This theme arose among participants who used security cameras or other security systems, or participants who mentioned concerns about door locks, which control physical access to the home.

> Most of the switches and bulbs are used to control the lightning in the home, for security purposes... The cameras are used for security, to be able to monitor the doors when we are not home, and the dog when we are away from home briefly. (P3)

Adversary	Concerned	Mentioned but not concerned
Unspecified bad actors	9/15	0/15
Company	1/15	8/15
Government	2/15	2/15
Owner of smart home devices	1/15	1/15
3rd party automation authors	1/15	0/15
No identified adversaries	2/15	

Table 4: Adversaries identified by participants.

Concern about physical security is perhaps natural, as devices like door locks and security cameras exist expressly for that purpose. Indeed, one participant cited these concerns as a reason to be more concerned about security risks with a smart home than with a laptop or phone:

> For the home, it's definitely something that I'm worried about because I don't want someone accessing the lock, or knowing the motion sensor data, or when we're home. On your phone, you have a degree of control — you can encrypt your phone, you can set up proper security — PIN and locks and stuff... I'm more concerned about my home than my phone, because if I lose my phone I can remotely wipe it. (P6)

Meanwhile, other risks with smart home devices occur as a side effect, such as privacy violations. Many participants acknowledged privacy could be an asset, particularly in the form of audio or behavior logs. However, half of these participants were not especially concerned about privacy risks:

> It's not like I openly admit to anything ridiculous that would incriminate me. And even if I did, no one's going to hear it, because Amazon doesn't release audio logs... That doesn't bother me, I guess — some people, it freaks them out, but it's not a big deal. It's just part of big data. They're just trying to gather data for advertising purposes, whatever floats their boat. (P5)

Other, less commonly identified assets that might be affected by security or privacy risks included bandwidth, money, or personally-identifiable information (PII). In one of our exploratory interviews, we heard an anecdote in which someone set up a custom smart sprinkler system which, due to an incorrect trigger, accidentally watered the lawn for a week and led to a significant water bill. Though this case was accidental, it could also be a compelling target for an attack.

Notably, no participant identified availability of device functionality as an asset that might be attacked. Although several participants voiced concerns about reliability (see Section 5.7), none connected this concern to security risks (rather identifying non-malicious network or power failures).

Adversaries. In general, when participants speculated about potential attacks on their smart home, they did not articulate specific adversaries in those scenarios (often referring to adversaries as "someone").

The most frequently identified potential adversaries were the companies that manufactured their smart home devices and that received data from those devices in the cloud. However, almost all participants who acknowledged this sort of behavior from companies were not concerned, and trusted the companies to protect their privacy. For example:

> In terms of the smart home stuff in particular, we are dealing with Amazon, we are dealing with big companies that are probably not totally irresponsible about privacy and security. (P11)

A few other participants mentioned the government as an adversary. However, they seemed to consider this concern only in the abstract sense, not providing many specifics on actions the government would take. For example, only one of the participants who mentioned the government also mentioned the murder case where law enforcement is requesting that Amazon turn over recorded audio data from an Echo device [3]. Less specifically, participants voiced general concerns about the government's surveillance capabilities and the current political climate (circa February 2017). For example:

> I am beefing up operational security in a big way, because I have spoke publicly against fascism, and I work in a publicly funded institution, I expected to be targeted at some point. (P1)

Other participants were aware of the government's potential surveillance capabilities but not overly concerned:

> I haven't changed any of my behavior in the house. If the FBI/CIA actually ever gets a recording of what's going into my Echo, they'll probably just think I'm a weirdo. (P8)

Participants had few concerns about the developers of smart home applications or custom automations as adversaries. P12 noted that the custom automations for the Vera hub were simple enough that he could read and understand it.

> Oh no, [the code] is so plain language. The only code they're writing for me is conditional commands. To turn on all the lights, I do that all myself, that's a standard scene... It's just the two tier deep programming [sic] that I've gotten their help with. And it's pretty obvious, the code they've written, I've saved it in a text file, it's you know, less than 30 characters. It's pretty obvious it's only pointing — it's like COBOL. (P12)

This lack of concern represents a gap to the threat models of security experts, who often explicitly include app developers as potential adversaries in their threat models and attempt to curtail the default capabilities of applications (e.g., [27]).

Finally, some participants were concerned or encountered issues with other residents in or visitors to the home; we discuss these issues in Section 5.6 below.

Vulnerabilities. Participants identified few concrete vulnerabilities that might lead to a security or privacy compromise, and no potential vulnerability was mentioned by a majority of participants (see Table 5). In general, we found that participants with different levels of technical knowledge identified different types of vulnerabilities in the threat model. For example, only participants with a more technically accurate mental model mentioned lack of transport level security (HTTPS) as a vulnerability.

Vulnerabilities	Concerned	Mentioned but not concerned
Data at risk in the cloud	1/15	5/15
Weak passwords	5/15	0/15
Lack of transport level security	4/15	0/15
Insecure devices	4/15	0/15
Malicious devices	3/15	0/15
Unsecured Wi-Fi network	2/15	0/15
Devices can be unpaired	1/15	0/15
No identified vulnerabilities	3/15	

Table 5: Vulnerabilities identified by participants.

> A lot of stuff is just totally unencrypted. Some of it is encrypted, a lot of it doesn't validate SSL certs. ... Even today, there's a lot of use of weak encryption ciphers. Yeah, it's pretty awful. (P1)

Meanwhile, participants with a less sophisticated mental model were more concerned about weak passwords and unsecured Wi-Fi networks, which are vulnerabilities that are not specific to the smart home context.

> People are concerned that someone could check into their camera or their lights... I guess they're not smart enough to know that they can't do that if they don't get your password. (P3)

Some participants mentioned concerns about malicious or vulnerable devices, either specifically (e.g., P8 was aware of Foscam web cam vulnerabilities [18]) or more generically:

> I think the biggest thing is just the amount of questionable things that have happened within the IoT space from some of the up and coming companies. That has me questioning what they can and can't do... I've just heard horror stories from some of the smaller companies. (P2)

Threats. As with vulnerabilities, there was not a particular threat or attack that a majority of participants were concerned about (see Table 6). While many acknowledged that companies or other adversaries *could* record and store private data, like audio/video feeds and behavioral logs, again we found that most were not concerned about it.

Again, we saw that participants with more advanced mental models voiced more concrete and technical threats, such as network attacks and network mapping. For example, P10 identified a specific threat: that an adversary with physical access to the home could un-pair a device from the user's hub, and re-pair it with their own hub.

Reasons for Lack of Concern. Even when participants were aware of security and privacy issues, they were often not actively concerned about them, voicing several reasons.

One reason for lack of concern, discussed above, is explicit trust in companies handling user data, such as Amazon.

Some participants were not concerned about attacks because they did not consider themselves a worthwhile target (notably, not considering untargeted attacks like widespread DDoS):

> I read some stuff about Hue bulbs being hacked, but I live in a small town. No one is going to pull up to my house and do any of that stuff. (P7)

Threats	Concerned	Mentioned but not concerned
Continuous audio/video recording	3/15	5/15
Data collection and mining	1/15	5/15
Adversarial remote control	4/15	1/15
Network attack on local devices	3/15	1/15
Spying by other user in home	3/15	0/15
Account/password hacking	2/15	0/15
Network mapping by mal. devices	1/15	0/15
Re-pair device with attacker's hub	1/15	0/15
No identified threats	1/15	

Table 6: Threats identified by participants.

Some believed they have nothing to hide, a perception that other researchers have reported for online behavior [15]. Others believed that they had taken sufficient steps to secure their systems, such as with strong passwords, so they did not need to worry further about security. For example:

> I also know many, many people who have such powerfully weak passwords, that if someone were driving around trying to get into someone's stuff, they would get into someone's stuff with weak passwords, and not into mine. (P3)

> I see the ability for devices to be manipulated if not secured properly, but from what I've read it seems like you can lock your system down pretty well, by just having a secure network and backup options. (P4)

Seven participants explicitly identified a tradeoff, requiring that one accepts security or privacy risks in exchange for the functionality and convenience of a smart home. For example:

> ...your data's going somewhere, and it comes down to who you are going to trust with it. You can trust it with Amazon, who has a record of everything you have spoken to your Echo, or are you gonna trust it with Google, who has access to your email, your map search history, your web search history? It depends on who you think is gonna do what with your data. ... It's a tradeoff of these free services — you're getting Gmail for free, but you're letting them run ads. (P6)

> I think our security is so compromised in so many different ways and I'm broadly speaking willing to accept some of the benefits of having these system understand my life — targeted advertising and various other conveniences. (P11)

5.5 Mitigation Strategies

Here, we consider approaches participants took to mitigate their security and privacy concerns. Mitigation strategies varied greatly, with no single strategy shared by more than five participants, suggesting that best practices for end user smart home security have not become standard.

Technical Mitigations. Two participants intentionally kept their smart home devices on a separate Wi-Fi network from other home electronics, perhaps concerned about attacks

by compromised smart home devices on other electronics, which may have more valuable data. These participants also blocked certain traffic from their devices: P1 blocked all unencrypted traffic, and P2 prevented their SmartThings hub from communicating with cloud servers, instead using an MQTT broker to control it from a local server.

Some participants attempted to mitigate password and Wi-Fi security related concerns with best practices, presumably learned from more traditional computing contexts:

> I don't have any security concerns because I feel fairly confident that the — I know that my passwords for all those accounts are very secure. (P3)

A few participants, with more technical backgrounds, desired additional security or privacy features on their devices, such as better use of HTTPS, or more granular permissions on the sensors on devices. P1 in particular wanted to be able to switch off the microphone on a Nest thermostat device. Only two participants mentioned deleting camera recordings or other logs of behavior to protect their privacy.

In some cases, participants used mitigation strategies with unclear benefits, suggesting limited underlying technical knowledge. For example, P7 only used Z-Wave smart home devices for security reasons, and when asked why, he said:

> I don't really remember. There was an article I was reading about... it was when I started out like two years ago that I researched it and I got these things in my head... I don't remember the specifics. I'm not an expert on any of this stuff. I try to do my research, but I have to take other people's opinion at face value. (P7)

Non-Technical Mitigations. A possible strategy for mitigating privacy risks in smart homes is simply altering one's behavior around those devices. For example, one might avoid saying certain things around the Echo or doing certain things in front of cameras. However, when asked explicitly about such behavior changes, nine participants explicitly mentioned that they did not change their behavior at all. Others only mentioned changing their behavior in theory:

> If I was to do something illegal I wouldn't do it in the room that has the Alexa and the camera in it. I would probably also turn off my cell phone, because... you don't know. I generally don't feel concerned because I'm not currently up to anything that is so private that it can't be stored in Amazon's temporary voice audio recording database. (P13)

Several participants made choices about where to place devices, or when those devices were enabled, for privacy reasons:

> We choose not to face [the camera towards] any interior portions. I do have a camera that's easy to set up, and when we go out of town for a couple of days, I'll just plug it in and it faces the interior, but never when we're actually home. (P1)

> With the camera I have in the house, I do have it plugged it into [a smart] powerstrip. So I don't really need that on when I'm there. So that's one thing that I guess we did do something little different, just have the camera come on when we're away. (P7)

5.6 Multi-User Interactions

We now turn to concerns and issues related to incidental users of the smart home, who were not primarily involved in selecting or automating devices. Three participants were incidental users, and we also asked primary users about disagreements with or concerns of incidental users.

Differences in Mental and Threat Models. We found that in general, incidental users of smart homes have simpler mental models, less awareness of security/privacy issues, and weaker threat models. This is perhaps natural; the person who wants to set up a smart home is likely more enthusiastic and curious about researching the technology, while the other resident(s) might simply tolerate their smart home "hobby".

For example, P14 lives in a fairly complex smart home set up by her husband (who is seemingly tech-savvy, as their OpenHAB hub requires programming skills). However, P14's mental model of their smart home is incomplete (see her drawing in Figure 1b), and she deferred most of the worrying about security to her husband. When asked specifically if they had security or privacy concerns, she said:

> It is something we joke about, but he's assured me that no one's going to be able to hack into it. I don't know if I believe that. (P14)

Differences in Access. Additionally, we found that incidental users do not always have full access to the smart home. Often they do not have the proper apps installed to control the home, either because the devices can be controlled without the app, using an Amazon Echo or Google Home, or these users were simply not interested in playing with the app and setting up automations on their own.

Differences in Power and Control. One consequence of non-primary users having less access and less interest in smart homes is that it leads to situations where the primary user may have — intentionally or unintentionally — more power over the other residents of the home. For example, we observed three such cases in our interviews.

Case 1: Restricted Access. P5 did not give their spouse access to the thermostat, because they wanted to keep it at a certain temperature to save power:

> I locked down my thermostat from [my wife] specifically. Because she complains that it is hot all the time, and I'm like, "Just turn on the fan, just turn on the ceiling fan and stand under it, and you'll be good," because it costs money. (P5)

Case 2: Audio/Video Surveillance. P13 lives in a house where the smart home setup was provided by the landlord. In particular, they had an Amazon Echo, a Nest surveillance camera, and Philips Hue lights. The landlord, being the owner of the devices, had accounts associated with these devices. That gave the landlord access to transcriptions and recordings of voice queries to the Echo, and could receive notifications from the security camera. The landlord accessed private data in at least one instance:

> We threw a party and didn't tell the woman who coordinates our house, and someone unplugged the Nest camera in the kitchen because they wanted to recharge their phone... and when it is unplugged, it automatically sends an email to whoever's account is associated with the camera, and it has a photo of the last thing the camera saw. So we were throwing this huge party, and it sent a photo of the kitchen..., and so the coordinator got the email and it was like "Your camera was unplugged, this is the last thing the camera saw!"... She wasn't mad! They were excited that we were having a party. (P13)

In this instance, there were no negative consequences, nor was P13 particularly concerned about their landlord's access to the smart home, but other situations may not be as benign.

Case 3: Behavioral Surveillance. P2 has an extensive smart home setup, and mentioned using the smart lock to find out when his wife and children arrived at home. In addition, he had custom software to detect when devices were on the network, which also indicates who is at home. When asked about disagreements with his wife, P2 said:

> My wife hates the aspect that I know when her device comes or goes on the local LAN, which obviously creates an audit log, so to speak, of when she's at home. She's now chiming in, that's the reason her phone doesn't connect to the Wi-Fi anymore, so I can't track her. (P2)

In this case, P2's surveillance does not appear to have been malicious but rather a result of his experimentation with the smart home — but again, other situations may be more dangerous (e.g., domestic abuse [36]).

Trolling. On a more lighthearted note, participants identified several instances of "trolling" among residents or guests in a smart home. Though these examples are not malicious and were not poorly received, they also highlight potential tensions that may arise between multiple users. For example:

> I had my family here over the weekend, and they have an Echo as well... They said "Hey Alexa, put poop on my shopping list" and then they said "Hey Alexa, order that", and of course it said "Are you sure?" and they let me say no. (P5)

5.7 Non-Security and Privacy Concerns

Finally, participants often cited non-security and privacy related concerns that influenced how they set up their smart home system. These concerns can be at odds with security and privacy, and researchers or platform designers focused on addressing security and privacy issues must consider these other constraints as well.

Reliability. Eight participants expressed concern about their home's resilience to network and power failures. In the event that their devices and hubs could no longer connect to the internet, participants wanted their devices to continue to function as normal, including their automations. For these participants, the ability to run automations locally was a deciding factor on which hub they decided to buy.

> If... you're trying to do something, and it doesn't work because the internet is down, that's really annoying... your wall switches should still work, your automations might not work, but simple stuff that doesn't require the internet to process things should still work. (P6)

As discussed above, despite this concern about reliability, no participant considered it in the context of security, i.e., no one mentioned that availability could be impacted maliciously. We also observe that maliciously induced failures could be leveraged for other attacks, e.g., to access door locks, although no participant voiced such a concern either.

Interoperability. Six participants mentioned that they want their devices to be interoperable, i.e., compatible with the rest of their smart home system. These participants would like their devices to be controllable by their hub, by their Echo/Google Home, and/or by a single, centralized app. As discussed above, in several cases participants installed third-party applications, such as IFTTT or Stringify, for the sole purpose of making devices interoperable. Such ad hoc connections potentially introduce new security vulnerabilites by expanding the attack surface of their system.

Cost. For some participants, a more prominent barrier to adoption was device cost. For example, this led P10 to cobble together his own "smart" sprinkler system rather than buying an existing smart device, increasing the risk of user error and potentially opening the door for security vulnerabilities.

5.8 Results Summary

Before stepping back in Section 6, we summarize our key findings from interviews with smart home end users:

- Participants have varied and sparse threat models, and do not share a common set of concerns or mitigations.
- Participants' threat models often depend on the sophistication of their technical mental models.
- Reasons for lack of concern about security/privacy issues include not feeling personally targeted, trusting potentially adversarial actors (like companies or governments), and believing their existing mitigation strategies to be sufficient.
- Concerns of security experts about smart homes, such as insecure or malicious devices, company data collection, attacks on device availability, or malicious or buggy third-party apps, were generally not shared by participants.
- Homes with multiple users pose unique security and privacy challenges, especially when the primary user has greater knowledge and control of the system than incidental users.
- Participants make smart home technology choices based on requirements that may conflict with security and privacy, including cost and interoperability.

6. DISCUSSION

We now step back to reflect on lessons from our findings, develop recommendation, and discuss study limitations.

6.1 Lessons

Incomplete mental models lead to gaps in threat models and security behaviors. Echoing prior work on mental models and security [34, 49], we found that participants with more sophisticated mental models had more advanced threat models that identified risks unique to smart

homes, and were able to take specific precautions to address these risks, such as blocking unencrypted traffic from their smart home devices. On the other hand, participants with less sophisticated mental models did not identify smart home-specific vulnerabilities and threats, and often based their mitigation strategies on best practices from other technologies, like using strong passwords, or adopted ad hoc strategies with unclear benefits, like avoiding using non-Zigbee devices.

The absence of common threat model elements and mitigation strategies suggests that best practices for smart home security have yet to be developed. This gap makes it difficult for users without a deep technical understanding of the technology to make informed security decisions.

Participants were more about physical security issues than privacy issues. The physical security of the home was a common concern voiced by our participants. This concern was expressed in two ways: either concern about attackers compromising security-critical devices like smart locks, or using the smart home to enhance their home security, with light timers and cameras. Brush et al. [9] found similar concerns: remote access to locks and cameras is important but creates a security risk.

However, most participants were unconcerned about privacy issues with their smart homes, despite having at least a cursory awareness. A possible explanation for this result could be that devices like door locks have security as their *primary* purpose, so a security failure would be equivalent to a functionality failure. By contrast, privacy risks with other devices are side-effects of their intended purpose, (e.g., privacy risks due to the Echo's ability to record audio).

This result could also be explained in part by our participant group: smart home users. These users have already chosen to set up a smart home (or had one set up for them); we did not hear from people who chose *not* to install a smart home due to security and privacy concerns. We discuss this limitation further in Section 6.3 below.

Mismatch between awareness and power of smart home administrator and other residents. In addition to replicating findings about the primary/incidental user dynamic from previous studies of end users of smart homes [9, 20, 41, 40], which found that most households have one user who is more active about researching, purchasing, and setting up smart home devices, our findings suggest that incidental users of smart homes may be less tech-savvy and/or less informed or aware about potential security and privacy issues. These discrepancies can lead to a power imbalance in a home, with the primary user in a position to (maliciously or not) spy on other residents or limit their control of the home.

A key observation here is that while the people who set up smart homes, particularly early adopters, often treat the technology as a personal hobby, smart homes are fundamentally *not* personal technologies. As a result, any security and privacy (or other) decisions made by the primary user directly affects other residents and visitor. If the primary and incidental users share a threat model, this interaction can be positive; however, if they do not agree on concerns or, worse, the primary user is adversarial (e.g., abusive) towards the incidental users, dangerous situations can arise.

Flexible end user programming limits usefulness of third-party applications. We found that users make limited use of third-party apps, e.g., on Samsung SmartThings. Instead, they more frequently use end-user programming interfaces (e.g., to set custom automation rules) or directly write scripts. When third-party apps were used, it was often to connect other ecosystems to the platform, e.g., to enable Amazon Echo based voice control of Samsung devices. This finding begs the question: Why? Are packaged apps not sufficiently flexible for diverse home environments (unlike on more homogenous smartphones)? Do they not yet provide sufficiently compelling functionality? Future work may shed light on these questions; in the meantime, the app platforms may not be the most critical place for the security community to focus its efforts — as otherwise seems natural, given the wealth of work on smartphone app platforms.

6.2 Recommendations and Future Work

We develop recommendations for the designers of smart home platforms and devices, as well as for future research.

UI/UX for User Awareness and Control. By improving users' technology mental models, we can also improve the accuracy of their threat models, enabling conscious decisions about whether to mitigate or ignore privacy or security risks. A possible strategy is to surface more information to users about what devices are doing — e.g., by providing usable auditing features in the associated phone apps, or by including physical indicators on devices (e.g., recording lights). Consolvo et al. used similar techniques for surfacing information leakage over unencrypted Wi-Fi networks [13]. Such indicators are already common for cameras and microphones, but may not be noticed by users performing unrelated tasks [46] — thus, future work must study how to design effective indicators in the smart home context, where users are often not directly interacting with devices.

Similarly, user control can be enhanced by ensuring that users can interact with devices physically, not only through apps. This can improve multi-user interactions and can help mitigate potential impacts of network failures, cloud outages, and phone or app problems. Indeed, several participants explicitly mentioned the need for physical switches.

> If the system is shut off, your wall switches should still work, your automations might not work, but simple stuff that doesn't require the internet to process things, it should still work. (P6)

Design Consciously for Multiple Users. In many of today's smart home platforms, support for multiple users is overlooked, and platform designers seem not yet to have deeply considered the potential risks among users in the same home. For example, users of SmartThings can easily monitor other users, and the Echo allows access to audio logs. From our interviews, we also heard about cases in which incidental users were intentionally or unintentionally denied access to smart home controls. As smart homes become more prevalent, similar issues may arise with guests.

Thus, future smart home platforms must take into account multi-user interactions and the potential power imbalance between the primary user and incidental (and often less tech-savvy) users. In addition to the need to support multiple distinct user accounts, usability and discoverability of

features are critical for secondary, less technical users.

The user control and awareness recommendations we make above can also help improve the multi-user experience. For example, if devices can all be controlled with physical switches, then all residents are guaranteed the ability to control that device. Similarly, physical recording indicators and other usable audit logs can help improve awareness of incidental users. We encourage future research to further study both the dynamics of multi-user smart homes as well as evaluate potential designs to mitigate these issues.

Reputation Systems for Smart Home Options. Not all users can (or should) become technical experts. Instead, external guidance may be required to help users make informed decisions about which products have stronger security and privacy properties. For smartphones, centralized app stores provide app reputation information, and prior work [25] has shown that users use these reviews to make decisions about which applications to install, including for security and privacy reasons. Similarly, the Electronic Frontier Foundation's secure messaging scorecard [23] aims to inform users about the security properties of different messaging options. The smart home ecosystem is much more heterogeneous, and there are no well-known centralized resources for security or privacy sensitive users to inform themselves. The recent news that Consumer Reports will begin evaluating products for security and privacy [14] is a step in the right direction.

Develop Standard Best Practices for End Users. As discussed above, we found that participants often adopted best security practices not specific to smart homes, such as strong password and Wi-Fi security. However, these practices do not cover many of the security and privacy issues unique to smart homes, and security experts must develop — and communicate effectively to end users — an updated set of best practices for smart home contexts. For example, one such recommendation might be unplugging or muting recording devices when they are not needed or during sensitive conversations, or alerting guests to their presence.

Design for Secure and Robust Interoperability. Interoperability between devices and smart home ecosystems (e.g., Amazon and SmartThings) was important to many participants, who often installed third-party apps or built custom solutions to connect different pieces of their smart home. Since security issues often arise at the boundaries between components, these user-created interoperability links (likely different across individual smart homes) may present future points of weakness. Security researchers should study these integrations, and smart home platform and devices designers should explicitly design for robust interoperability.

Minimize Tradeoffs for Security and Privacy. Many participants identified a tradeoff between security and privacy with functionality and convenience, in some cases sounding resigned to it. We challenge smart home designers and researchers to present a better tradeoff. For example, certain technical design choices can reduce risks without significantly impacting functionality, like not requiring the cloud to run automations. (Indeed, SmartThings initially only supported running apps in their cloud, but now supports apps on the local hub [51] — although perhaps for reliability rather than security reasons.) By minimizing these tradeoffs when possible, we can remove the decision-making burden from users and enable adoption of smart home technologies by people who are not willing to make the tradeoffs required today.

6.3 Limitations

Finally, we reflect on several limitations of this work. First, we only interviewed participants living in smart homes, not people who chose not to install them for security or privacy (or other) reasons. Future work should study this deliberate non-user population, as they may have more pronounced concerns that hindered adoption in the first place. Here, our focus was on participants who could speak to concrete, rather than hypothetical, smart home experiences.

Second, our sample skews towards primary users and smart home enthusiasts, despite our efforts to recruit more passive users of smart homes. This is likely in part due to self-selection bias among people drawn to participate in our study, and because we recruited from smart home-focused online communities. Nevertheless, these participants' accounts of other residents, as well as the interviews with three non-primary users, help shed light on this class of end users.

Third, smart home technology is new and still developing — commercial platforms targeted at non-technical consumers, like Samsung SmartThings or Amazon Echo, are recent developments. Thus, our participants are among the earliest adopters; they may be more willing to choose convenience over security or privacy, or be generally more tech-savvy, than non-adopters. As smart homes become more widespread, the makeup of the user base will shift, and future work should consider these changes. Meanwhile, our findings already shed light on issues that will arise and become more complex, e.g., around multi-user scenarios, as adoption increases.

Finally, this qualitative study, by its nature, does not produce quantitative conclusions, e.g., about the prevalence of certain concerns or lack thereof. This work lays the groundwork for future quantitative studies to investigate such questions.

7. CONCLUSION

Consumer smart home technologies are becoming increasingly prevalent. Alongside the convenience offered by these technologies, they raise new security and privacy risks. Though researchers have begun studying these technologies themselves, there has been little study of the end users of modern smart homes: what are their mental models, security and privacy concerns, mitigation strategies, and how does the presence of multiple users compound these issues? We sought to answer these questions in our work, conducting in-depth interviews with fifteen participants (twelve smart home administrators and three other residents). Our findings shed light on their mental models and security concerns (or lack thereof) — for example, revealing incomplete threat models and ad hoc mitigation strategies based on best practices for older technologies — and highlight potential tensions between multiple smart home users. These findings lay the groundwork for continued study of smart home end users as the technologies develop further and see increased adoption, and we provide recommendations to smart home technology designers and researchers for where to focus future efforts. For example, we highlight the need to help shape user mental models, consciously design for multiple users, and design for security and privacy alongside key features valued by users (e.g., interoperability and remote access).

Acknowledgements

We are especially grateful to our user study participants, as well as our pilot study participants, Camille Cobb and Alex Takakuwa. We thank our anonymous reviewers and our shepherd, Blase Ur, for helpful feedback on an earlier version. We also thank Luis Ceze for useful conversations about smart home setups, as well as Yoshi Kohno, Kiron Lebeck, Lucy Simko, and Anna Kornfeld Simpson for reviewing an earlier draft. This work was supported in part by the National Science Foundation under Awards CNS-1513584 and CNS-1565252, and by a Hacherl Endowed Fellowship.

8. REFERENCES

[1] Vera smarter home control. Accessed March 7, 2017, Online at http://getvera.com/.

[2] Wink hub. Accessed March 7, 2017, Online at https://www.wink.com/products/wink-hub/.

[3] Amazon shares data with Arkansas prosecutor in murder case, Mar. 2017. Accessed March 6, 2017, Online at http://bigstory.ap.org/article/1110e4449c3f4191909e4010da935056/amazon-shares-data-arkansas-prosecutor-murder-case.

[4] Amazon Echo. Accessed March 7, 2017, Online at https://www.amazon.com/Amazon-Echo-Bluetooth-Speaker-with-WiFi-Alexa/dp/B00X4WHP5E.

[5] Apple HomeKit. Accessed March 7, 2017, Online at https://www.apple.com/ios/home/.

[6] O. Arias, J. Wurm, K. Hoang, and Y. Jin. Privacy and security in Internet of Things and wearable devices. *IEEE Transactions on Multi-Scale Computing Systems*, 1(2):99–109, Nov. 2016. DOI 10.1109/TMSCS.2015.2498605.

[7] K. Baxter, C. Courage, and K. Caine. *Understanding Your Users: A Practical Guide to User Research Methods*. Morgan Kaufmann, second edition, 2015.

[8] N. Bilton. Nest thermostat glitch leaves users in the cold. The New York Times, Jan. 2016. Accessed March 7, 2017, Online at https://www.nytimes.com/2016/01/14/fashion/nest-thermostat-glitch-battery-dies-software-freeze.html.

[9] A. J. Brush, B. Lee, R. Mahajan, and S. Agarwal. Home automation in the wild: Challenges and opportunities. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2011. DOI 10.1145/1978942.1979249.

[10] J. bum Woo and Y. kyung Lim. User experience in do-it-yourself-style smart homes. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2015. DOI 10.1145/2750858.2806063.

[11] D. J. Butler, J. Huang, F. Roesner, and M. Cakmak. The privacy-utility tradeoff for remotely teleoperated robots. In *Proceedings of the Annual ACM/IEEE International Conference on Human-Robot Interaction*, 2015. DOI 10.1145/2696454.2696484.

[12] E. K. Choe, S. Consolvo, J. Jung, B. L. Harrison, S. N. Patel, and J. A. Kientz. Investigating receptiveness to sensing and inference in the home using sensor proxies. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2012. DOI 10.1145/2370216.2370226.

[13] S. Consolvo, J. Jung, B. Greenstein, P. Powledge, G. Maganis, and D. Avrahami. The Wi-Fi privacy ticker: improving awareness & control of personal information exposure on Wi-Fi. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2010. DOI 10.1145/1864349.1864398.

[14] Consumer Reports. Consumer Reports to begin evaluating products, services for privacy and data security. Accessed March 7, 2017, Online at http://www.consumerreports.org/privacy/consumer-reports-to-begin-evaluating-products-services-for-privacy-and-data-security/.

[15] G. J. Conti and E. Sobiesk. An honest man has nothing to fear: User perceptions on web-based information disclosure. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2007. DOI 10.1145/1280680.1280695.

[16] K. L. Courtney, G. Demeris, M. Rantz, and M. Skubic. Needing smart home technologies: The perspectives of older adults in continuing care retirement communities. *Informatics in Primary Care*, 16(3):195–201, 2008.

[17] A. Cui and S. J. Stolfo. A quantitative analysis of the insecurity of embedded network devices: Results of a wide-area scan. In *Proceedings of the Annual Computer Security Applications Conference*, pages 97–106. ACM, 2010. DOI 10.1145/1920261.1920276.

[18] CVE. CVE-2013-2560, 2013. Accessed March 7, 2017, Online at https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2013-2560.

[19] N. Dell, V. Vaidyanathan, I. Medhi, E. Cutrell, and W. Thies. "Yours is better!": Participant response bias in HCI. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2012. DOI 10.1145/2207676.2208589.

[20] A. Demeure, S. Caffiau, E. Elias, and C. Roux. Building and using home automation systems: A field study. In *Proceedings of International Symposium on End User Development (IS-EUD)*, 2015. DOI 10.1007/978-3-319-18425-8_9.

[21] T. Denning, T. Kohno, and H. M. Levy. Computer security and the modern home. *Communications of the ACM*, 56(1):94, Jan. 2013. DOI 10.1145/2398356.2398377.

[22] T. Denning, C. Matuszek, K. Koscher, J. R. Smith, and T. Kohno. A spotlight on security and privacy risks with future household robots: Attacks and lessons. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2009. DOI 10.1145/1620545.1620564.

[23] Electronic Frontier Foundation. Secure messaging scorecard. Accessed March 7, 2017, Online at https://www.eff.org/secure-messaging-scorecard.

[24] N. Feamster, S. Grover, and R. Ensafi. Who will secure the Internet of Things?, Jan. 2016. Accessed March 7, 2017, Online at https://freedom-to-tinker.com/2016/01/19/who-will-secure-the-internet-of-things/.

[25] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2012. DOI 10.1145/2335356.2335360.

[26] E. Fernandes, J. Jung, and A. Prakash. Security analysis of emerging smart home applications. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2016. DOI 10.1109/SP.2016.44.

[27] E. Fernandes, J. Paupore, A. Rahmati, D. Simionato, M. Conti, and A. Prakash. FlowFence - Practical data protection for emerging IoT application frameworks. In *Proceedings of the USENIX Security Symposium (USENIX Security)*, 2016.

[28] J. L. Fleiss, B. Levin, and M. C. Paik. *Statistical Methods for Rates and Proportions*. John Wiley & Sons, 3 edition, 2003.

[29] Google Home. Accessed March 7, 2017, Online at https://madeby.google.com/home/.

[30] J. Granjal, E. Monteiro, and J. Sa Silva. Security for the Internet of Things: A survey of existing protocols and open research issues. *IEEE Communications Surveys and Tutorials*, 17(3):1294–1312, 2015. DOI 10.1109/COMST.2015.2388550.

[31] K. Hill. 'Baby monitor hack' could happen to 40,000 other Foscam users, Aug. 2013. Accessed March 4, 2017, Online at https://www.forbes.com/sites/kashmirhill/2013/08/27/baby-monitor-hack-could-happen-to-40000-other-foscam-users/#40c00e3a58b5.

[32] G. Ho, D. Leung, P. Mishra, A. Hosseini, and D. Song. Smart locks: Lessons for securing commodity Internet of Things devices. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2016. DOI 10.1145/2897845.2897886.

[33] D. Jonassen and Y. H. Cho. Externalizing mental models with mindtools. In *Understanding models for learning and instruction*, pages 145–159. Springer, 2008.

[34] R. Kang, L. A. Dabbish, N. Fruchter, and S. B. Kiesler. "My data just goes everywhere:" User mental models of the internet and implications for privacy and security. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2015. Online at https://www.usenix.org/system/files/conference/soups2015/soups15-paper-kang.pdf.

[35] F. Kawsar and A. J. B. Brush. Home computing unplugged: Why, where and when people use different connected devices at home. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2013. DOI 10.1145/2493432.2493494.

[36] T. Matthews, K. O'Leary, A. Turner, M. Sleeper, J. P. Woelfer, M. Shelton, C. Manthorne, E. F. Churchill, and S. Consolvo. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2017. DOI 10.1145/3025453.3025875.

[37] M. L. Mazurek, J. P. Arsenault, J. Bresee, N. Gupta, I. Ion, C. Johns, D. Lee, Y. Liang, J. Olsen, B. Salmon, R. Shay, K. Vaniea, L. Bauer, L. F. Cranor, G. R. Ganger, and M. K. Reiter. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 645–654, 2010. DOI 10.1145/1753326.1753421.

[38] A. McLean. IoT malware and ransomware attacks on the incline: Intel Security, Sept. 2015. Accessed March 4, 2017, Online at http://www.zdnet.com/article/iot-malware-and-ransomware-attacks-on-the-incline-intel-security/.

[39] E. McReynolds, S. Hubbard, T. Lau, A. Saraf, M. Cakmak, and F. Roesner. Toys that listen: A study of parents, children, and internet-connected toys. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2017. DOI 10.1145/3025453.3025735.

[40] S. Mennicken, J. Hofer, A. K. Dey, and E. M. Huang. Casalendar: A temporal interface for automated homes. In *Proceedings of the Conference on Human Factors in Computing Systems: Extended Abstracts (CHI EA)*, 2014. DOI 10.1145/2559206.2581321.

[41] S. Mennicken and E. M. Huang. Hacking the natural habitat: An in-the-wild study of smart homes, their development, and the people who live in them. In *Proceedings of the International Conference on Pervasive Computing (Pervasive)*, 2012. DOI 10.1007/978-3-642-31205-2_10.

[42] P. Morgner, S. Mattejat, and Z. Benenson. All your bulbs are belong to us: Investigating the current state of security in connected lighting systems. *CoRR*, abs/1608.03732, 2016.

[43] L. H. Newman. The botnet that broke the Internet isn't going away. Wired, Dec. 2016. Accessed March 7, 2017, Online at https://www.wired.com/2016/12/botnet-broke-internet-isnt-going-away/.

[44] T. Oluwafemi, T. Kohno, S. Gupta, and S. Patel. Experimental security analyses of non-networked compact fluorescent lamps: A case study of home automation security. In *Proceedings of the Learning from Authoritative Security Experiment Results (LASER)*, 2013. Online at https://www.usenix.org/laser2013/program/oluwafemi.

[45] Philips Hue. Accessed March 7, 2017, Online at http://www2.meethue.com/en-US.

[46] R. S. Portnoff, L. N. Lee, S. Egelman, P. Mishra, D. Leung, and D. Wagner. Somebody's watching me?: Assessing the effectiveness of webcam indicator lights. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2015. DOI 10.1145/2702123.2702164.

[47] F. Raja, K. Hawkey, and K. Beznosov. Revealing hidden context: Improving mental models of personal firewall users. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2009.

[48] The home automation market by the numbers. Remotely, June 2015. Accessed March 7, 2017, Online at https://blog.remotely.com/2015/06/20/the-home-automation-market-by-the-numbers/.

[49] K. Renaud, M. Volkamer, and A. Renkema-Padmos. Why doesn't Jane protect her privacy? In *Proceedings of the International Symposium on Privacy Enhancing Technologies (PETS)*, 2014. DOI 10.1007/978-3-319-08506-7_13.

[50] E. Ronen, C. O'Flynn, A. Shamir, and A.-O. Weingarten. IoT goes nuclear: Creating a ZigBee chain reaction. Cryptology ePrint Archive, Report 2016/1047, 2016. Accessed March 7, 2017, Online at http://eprint.iacr.org/2016/1047.

[51] Samsung SmartThings. Local processing. Accessed

March 7, 2017, Online at https://support.smartthings.com/hc/en-us/articles/209979766-Local-processing.

[52] S. Shane, M. Mazzetti, and M. Rosenberg. WikiLeaks releases trove of alleged C.I.A. hacking documents. The New York Times, Mar. 2017. Accessed March 7, 2017, Online at https://www.nytimes.com/2017/03/07/world/europe/wikileaks-cia-hacking.html.

[53] S. Sicari, A. Rizzardi, L. A. Grieco, and A. Coen-Porisini. Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76:146–164, Jan. 2015. DOI 10.1016/j.comnet.2014.11.008.

[54] A. K. Simpson, S. N. Patel, F. Roesner, and T. Kohno. Securing vulnerable home IoT devices with an in-hub security manager. Technical Report UW-CSE-2017-01-01, University of Washington, 2017.

[55] SmartThings. Accessed March 7, 2017, Online at https://www.smartthings.com/.

[56] Smart home. Statista Digital Market Outlook. Accessed March 4, 2017, Online at https://www.statista.com/outlook/279/109/smart-home/united-states.

[57] D. Townsend, F. Knoefel, and R. Goubran. Privacy versus autonomy: a tradeoff model for smart home monitoring technologies. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 4749–4752. IEEE, 2011.

[58] K. Toyama. *Geek heresy: Rescuing social change from the cult of technology*. PublicAffairs, 2015.

[59] B. Ur, J. Jung, and S. Schechter. The current state of access control for smart devices in homes. In *Workshop on Home Usable Privacy and Security (HUPS)*, 2013.

[60] B. Ur, J. Jung, and S. E. Schechter. Intruders versus intrusiveness: Teens' and parents' perspectives on home-entryway surveillance. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2014. DOI 10.1145/2632048.2632107.

[61] P. Worthy, B. Matthews, and S. Viller. Trust me: Doubts and concerns living with the Internet of Things. In *Proceedings of the ACM Conference on Doubts and Concerns Living with the Internet of Things*, June 2016. DOI 10.1145/2901790.2901890.

[62] D. J. Wu, A. Taly, A. Shankar, and D. Boneh. Privacy, discovery, and authentication for the Internet of Things. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS)*, pages 301–319. Springer International Publishing, 2016. DOI 10.1007/978-3-319-45741-3_16.

[63] T. Yu, V. Sekar, S. Seshan, Y. Agarwal, and C. Xu. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the Internet-of-Things. In *Proceedings of the Workshop on Hot Topics in Networks (HotNets Workshop)*, 2015. DOI 10.1145/2834050.2834095.

APPENDIX

A. INTERVIEW PROTOCOL

General Questions:

1. What devices do you own?
2. Can you briefly describe what you use them for?
3. What apps or automations do you have installed?
4. Do you access devices remotely, or only when you're physically in your home?

Mental Models:
For this next part, I'd like you to draw a diagram of how all of your devices are connected together. I can either email you a link to a Google Docs drawing that we can both edit, or you can draw it on a piece of paper and send it to me.

Security and Privacy Concerns:
(Start with questions not explicitly about security or privacy:)

1. When setting up your home, did you have any hesitations about getting any devices?
2. Are there any devices you thought about getting but decided not to get? Why?
3. Are there any devices that you used to use but later deactivated?

(Move on to direct questions if they have not already started talking about security and privacy:)

1. One type of concern we're interested in is security or privacy concerns. Do you or did you have any concerns like that about your smart home? You might not have any such concerns – that's fine, and we'd like to hear about that too.

   (OR, if security/privacy have been brought up organically:)

   Do you or did you have any other security or privacy concerns that you haven't mentioned yet?
2. Have you heard about any security or privacy issues with smart homes in the news? If so, did that news concern you, or do you think those issues are a little overblown?
3. How would you compare your level of security/privacy concern about your smart home devices to your level of concern about your phone or laptop computer?
4. *(For Echo/security camera users:)* Do you ever look at the audio/video logs of your Echo/camera?

Mitigation Strategies:

1. Thinking specifically about security and privacy concerns, have those concerns caused you to change any of your behaviors?
   (a) For example, do you act differently in your home around your smart devices?
   (b) Do you do anything to your devices – such as muting them – to mitigate your security or privacy concerns?
2. What kind of policies or controls would you like to have in your smart home to alleviate your security and/or privacy concerns?

Multi-User Scenarios:

1. How many people live in your home?
2. Who has access to the smart home?
3. Have you ever had disagreements with people in your home about how your smart home is set up?
4. Does everyone who has access have the same level of access? *(If yes:)* Have you had situations where you wanted someone to have limited access, and if so, how did you handle that?
5. Have you ever have situations where houseguests have interacted with your smart home? Did anything go wrong? Did anyone voice any opinions or concerns?

Failures:
Are there any other things that have gone wrong while setting up or using your smart home devices that you'd like to share?

Self-Reporting Technical Skills:
On a scale of 1-5:

1. How familiar are you with technology in general?
2. How familiar are you with computer security?
3. How familiar are you with smart home technology?

Closing Questions:

1. Are there any questions you expected me to ask?
2. Is there anything else you want to tell me about your smart home?