



Rethinking Access Control and Authentication for the Home Internet of Things (IoT)

Weijia He, *University of Chicago*; Maximilian Golla, *Ruhr-University Bochum*;
Roshni Padhi and Jordan Ofek, *University of Chicago*; Markus Dürmuth, *Ruhr-University
Bochum*; Earlene Fernandes, *University of Washington*; Blase Ur, *University of Chicago*

<https://www.usenix.org/conference/usenixsecurity18/presentation/he>

**This paper is included in the Proceedings of the
27th USENIX Security Symposium.**

August 15–17, 2018 • Baltimore, MD, USA

ISBN 978-1-931971-46-1

**Open access to the Proceedings of the
27th USENIX Security Symposium
is sponsored by USENIX.**

Rethinking Access Control and Authentication for the Home Internet of Things (IoT)

Weijia He, Maximilian Golla[†], Roshni Padhi, Jordan Ofek,
Markus Dürmuth[†], Earlence Fernandes[‡], Blase Ur
University of Chicago, [†] Ruhr-University Bochum, [‡] University of Washington

Abstract

Computing is transitioning from single-user devices to the Internet of Things (IoT), in which multiple users with complex social relationships interact with a single device. Currently deployed techniques fail to provide usable access-control specification or authentication in such settings. In this paper, we begin reenvisioning access control and authentication for the home IoT. We propose that access control focus on IoT capabilities (i. e., certain actions that devices can perform), rather than on a per-device granularity. In a 425-participant online user study, we find stark differences in participants' desired access-control policies for different capabilities within a single device, as well as based on who is trying to use that capability. From these desired policies, we identify likely candidates for default policies. We also pinpoint necessary primitives for specifying more complex, yet desired, access-control policies. These primitives range from the time of day to the current location of users. Finally, we discuss the degree to which different authentication methods potentially support desired policies.

1 Introduction

Recent years have seen a proliferation of Internet of Things (IoT) devices intended for consumers' homes, including Samsung SmartThings [35], the Amazon Echo voice assistant [2], the Nest Thermostat [48], Belkin's Wemo devices [5], and Philips Hue lights [32]. To date, IoT security and privacy research has focused on such devices' insecure software-engineering practices [3,13,15], improper information flows [15,40,45], and the inherent difficulties of patching networked devices [49,51].

Surprisingly little attention has been paid to *access-control-policy specification* (expressing which particular users, in which contexts, are permitted to access a resource) or *authentication* (verifying that users are who they claim to be) in the home IoT. This state of af-

fairs is troubling because the characteristics that make the IoT distinct from prior computing domains necessitate a rethinking of access control and authentication. Traditional devices like computers, phones, tablets, and smart watches are generally used by only a single person. Therefore, once a user authenticates to their own device, minimal further access control is needed. These devices have screens and keyboards, so the process of authentication often involves passwords, PINs, fingerprint biometrics, or similar approaches [6].

Home IoT devices are fundamentally different. First, numerous users interact with a single home IoT device, such as a household's shared voice assistant or Internet-connected door lock. Widely deployed techniques for specifying access-control policies and authenticating users fall short when multiple users share a device [50]. Complicating matters, users in a household often have complex social relationships with each other, changing the threat model. For example, mischievous children [38], parents curious about what their teenagers are doing [44], and abusive romantic partners [29] are all localized threats amplified in home IoT environments.

Furthermore, few IoT devices have screens or keyboards [37], so users cannot just type a password. While users could possibly use their phone as a central authentication mechanism, this would lose IoT devices' hands-free convenience, while naïve solutions like speaking a password to a voice assistant are often insecure.

Real-world examples of the shortcomings of current access-control-policy specification and authentication for home IoT devices have begun to appear. A Burger King TV commercial triggered Google Home voice assistants to read Wikipedia pages about the Whopper [47], while the cartoon South Park mischievously triggered Amazon Echo voice assistants to fill viewers' Amazon shopping carts with risqué items [34]. While these examples were relatively harmless, one could imagine a rogue child remotely controlling the devices in a sibling's room to annoy them, a curious babysitter with temporary

access to a home perusing a device’s history of interactions, or an enterprising burglar asking a voice assistant through a cracked window to unlock the front door [42].

In this paper, we take a first step toward rethinking the specification of access-control policies and authentication for the home IoT. We structure our investigation around four research questions, which we examine in a 425-participant user study. These research questions are motivated by our observation that many home IoT devices combine varied functionality in a single device. For example, a home hub or a voice assistant can perform tasks ranging from turning on the lights to controlling the door locks. Current access control and authentication is often based on a device-centric model where access is granted or denied per device. We move to a capability-centric model, where we define a capability as a particular action (e. g., ordering an item online) that can be performed on a particular device (e. g., a voice assistant). Intuition suggests that different capabilities have different sensitivities, leading to our first research question:

RQ1: Do desired access-control policies differ among capabilities of single home IoT devices? (Section 6.2 and 6.3).

We investigated this question by having each study participant specify their desired access-control policy for one of 22 home IoT capabilities we identified. For household members of six different relationships (e. g., spouse, child, babysitter), the participant specified when that person should be allowed to use that capability. Our findings validated our intuition that policies about capabilities, rather than devices, better capture users’ preferences. Different capabilities for voice assistants and doors particularly elicited strikingly different policies.

While the ability to specify granularly who should be able to use which capabilities is necessary to capture users’ policies, it incurs a steep usability cost. To minimize this burden through default policies, we asked:

RQ2: For which pairs of relationships (e. g., child) and capabilities (e. g., turn on lights) are desired access-control policies consistent across participants? These can be default settings (Section 6.4).

In our study, nearly all participants always wanted their spouses to be able to use capabilities other than log deletion at all times. Participants also wanted others to be able to control the lights and thermostat while at home. As intimated by the prior policy, the context in which a particular individual would use a capability may also matter. Children might be permitted to control lights, but perhaps not to turn the lights on and off hundreds of times in succession as children are wont to do. Nor should children be permitted to operate most household devices when they are away from home, particularly devices in siblings’ rooms. A babysitter unlocking the door from inside the house has far fewer security implications

than the babysitter setting a persistent rule to unlock the front door whenever anyone rings the doorbell.

RQ3: On what *contextual factors* (e. g., location) do access-control policies depend? (Section 6.5).

In addition to a user’s location, we found that participants wanted to specify access-control policies based on a user’s age, the location of a device, and other factors. Almost none of these contextual factors are supported by current devices. Finally, to identify promising directions for designing authentication mechanisms in the home IoT, we asked:

RQ4: What types of authentication methods balance convenience and security, holding the potential to successfully balance the consequences of falsely allowing and denying access? (Section 6.6).

Analyzing consequences participants noted for falsely allowing or denying access to capabilities, we identify a spectrum of methods that seem promising for authenticating users (Section 7), thereby enabling enforcement of users’ desired access-control policies for the home IoT.

Contributions We begin to reenvision access control and authentication for the home IoT through a 425-participant user study. Our contributions include:

- (i) Proposing access-control specification for the multi-user home IoT based on capabilities that better fits users’ expectations than current approaches.
- (ii) Showing the frequent context-dependence of access-control policies, identifying numerous contextual factors that future interfaces should support.
- (iii) Setting an agenda for authentication in the home IoT based on methods that minimize the consequences of falsely allowing or denying access.

2 Background

In this section, we scope our notion of home IoT devices, identify our threat model, and review current devices’ support for access control and authentication. We define home IoT devices to be small appliances that are Internet-connected and used primarily in the home. Internet-connected lights and thermostats are two examples. Many such devices are managed through a hub that facilitates communication between devices, enforces policies, and often allows for the creation of end-user programs or the use of apps.

2.1 Threat Model

The two major classes of adversaries in the smart home are external third parties and those who have legitimate physical access to the home. The former class includes those who exploit software vulnerabilities in

platforms [13], devices [3] (e.g., with Mirai), or protocols [16] intending to cause physical, financial, or privacy-related damage. The latter class includes household members with legitimate digital or physical access to the home, such as temporary workers or children [38]. These insider threats have received far less research attention, but are the focus of this paper. Insiders might be motivated to subvert a smart-home system's access controls for reasons ranging from curiosity to willful disobedience (e.g., a child attempting to take actions forbidden by their parents), or to attempt to correct imbalances created by the introduction of devices whose surveillance implications grant asymmetric power to certain members of a household (e.g., a parent tracking a teenager [44]).

We assume a domestic setting where occupants control home IoT devices through smartphones, voice assistants, rules, and physical interaction. For example, a maintenance worker may unlock the front door using a smartphone app, while a child might turn off their lights by speaking to a voice assistant. We aim for access-control rules that balance security, privacy, and functionality.

2.2 Affordances of Current Devices

Current home IoT devices have relatively limited affordances for access control and authentication. Taking a five-year-old survey of the home IoT landscape as a starting point [43], we surveyed current devices' affordances; Figure 1 shows representative samples. To control many current devices, people use smartphone apps that must be paired with devices. These apps offer various access-control settings. For example, the Nest Thermostat supports a binary model where additional users either have full or no access to all of the thermostat's capabilities. The August Smart Lock offers a similar model with guest and owner levels. Withings wireless scales let users create separate accounts and thus isolate their weight measurements from other users. On Apple HomeKit, one can invite additional users, restricting them to: (a) full control, (b) view-only control, (c) local or remote control.

Some devices offer slightly richer access-control-policy specification. The Kwikset Kevo Smart Lock allows access-control rules to be time-based; an owner can grant access to a secondary user for a limited amount of time. We find in our user study that time is a desirable contextual factor, but one of only many. We focus on capabilities, rather than devices. While most current devices do not allow for access-control policies that distinguish by capability, Samsung SmartThings lets users restrict third-party apps from accessing certain capabilities [36]. We find that restricting users, not just apps, access to a particular capability is necessary.

From this analysis, we found current mechanisms to be rudimentary and lack the necessary vocabulary for

specifying access-control rules in complex, multi-user environments. We aim to establish a richer vocabulary.

Current authentication methods for the home IoT appear transplanted from smartphone and desktop paradigms. Passwords are widely used in conjunction with smartphones. For example, SmartThings has an app through which a user can control devices. A user first authenticates to this app using a password. Voice-based authentication is currently very rudimentary and is not used for security, but for personalization. For instance, Google Home uses speaker recognition for customizing reminders, but not for security-related tasks [19].

3 Related Work

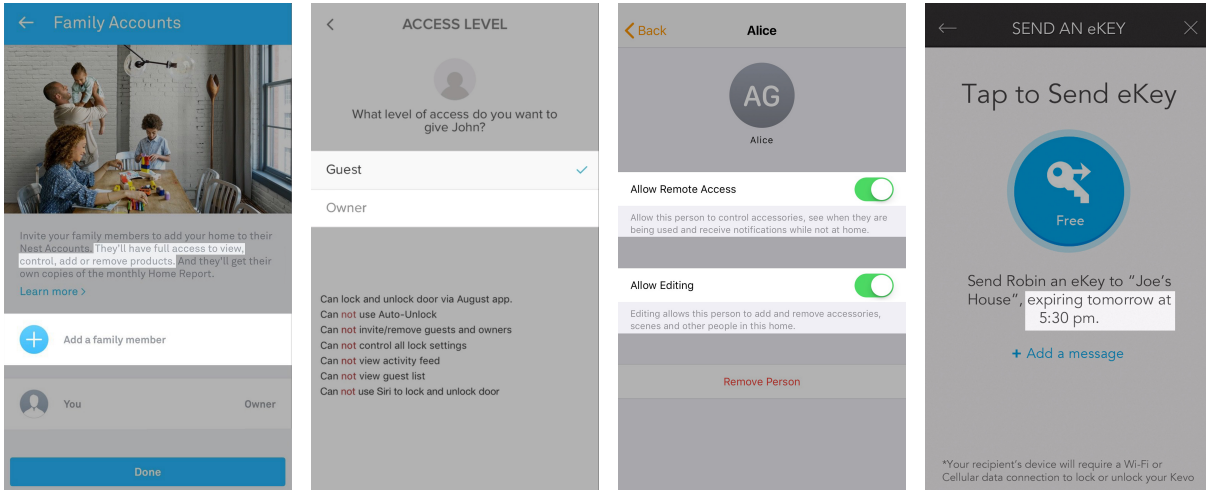
Current research focuses on analyzing and fixing the security of platforms [13, 14, 45], protocols [16], and devices [3]. Fernandes et al. discuss how smart-home apps can be overprivileged in terms of their access to devices and present attacks exploiting deficiencies in apps' access-control mechanisms [13]. Mitigations have involved rethinking permission granting [13, 22, 41].

Comparatively little work has focused on authorizing and authenticating *humans to home IoT devices*. Prior work has focused on the difficulties of access control in the home [4, 24, 25, 30], rather than solutions. Furthermore, the consumer device landscape has changed rapidly in the years since these initial studies.

Some older work has examined authentication [39] and access-control [43] for deployed home IoT devices, finding such affordances highly ineffective. Recent studies [31, 50] have sought to elicit users' broad security and privacy concerns with IoT environments, particularly noting multi-user complexity as a key security challenge. This complexity stems from the social ties in a home IoT setting. For instance, researchers have noted that roommates [26], guests [23], neighbors [7], and children [8, 38] are all important considerations in multi-user environments. We build on this work, identifying desired access-control rules for home IoT devices and bringing both relationships between home occupants and devices' individual capabilities to the forefront.

Prior research on IoT authentication has focused on protocols (e.g., Kerberos-like frameworks [1, 27]) without considering the constraints of users. Feng et al. introduced VAuth, voice-based authentication for voice assistants [12]. VAuth requires the use of wearable hardware to establish an authentication channel, however. One of our goals (RQ4) is to identify the authentication mechanisms that might be suitable for multi-user devices.

Smartphones can be considered a predecessor to the IoT, yet the large literature [9, 10, 11, 46] on specifying which apps can access which resources translates only partially to home IoT devices. Enck et al. discuss how



(a) Nest Learning Thermostat (b) August Smart Lock (c) Apple HomeKit (d) Kwikset Kevo Smart Lock

Figure 1: Current access-control-specification interfaces: The Nest Thermostat (a) only allows “all-or-nothing” specification, while the August Smart Lock (b) only offers coarse-grained access control via predefined Guest and Owner groups. In contrast, Apple’s HomeKit (c) differentiates between view and edit access level, as well as local and remote access. The Kwikset Kevo Smart Lock (d) provides time-based access control, but not other factors.

apps could gain access to resources by requesting permission from the user [9], while Felt et al. discuss how users may not always pay attention to such prompts [11]. A common theme is that apps access phone resources, and a phone is a single-user device not typically shared with others. On current versions of Android, one can configure secondary accounts with restrictions on what apps may be used [17], yet having separate accounts does not solve the multi-user challenges of home IoT devices.

4 Pre-Study

As a first step in exploring access control based on capabilities and relationships in the home IoT, we conducted a pre-study to identify capabilities and relationships that elicit representative or important user concerns. To ground our investigation of capabilities of the home IoT in devices consumers would likely encounter, we created a list of home IoT devices (Appendix A) from consumer recommendations in CNET, PCMag, and Tom’s Guide [33]. We grouped devices by their core functionality into categories including *smart-home hubs*, *door locks*, and *voice assistants*.

For each category of device, we collected the capabilities offered by currently marketed devices in that category. We added likely future capabilities, as well as the ability to write end-user programs [40, 45]. We showed each pre-study participant all capabilities identified for a single given class of device. The participant answered questions about the positive and negative consequences of using that capability, and they also identified addi-

tional capabilities they expected the device to have. We used this process to identify a comprehensive, yet diverse, set of capabilities that range from those that elicit substantial concerns to those that elicit none.

To identify a small set of relationships to investigate in the main study, we also showed participants a table of 24 relationships (e. g., teenage child, home health aide) and asked them to group these relationships into five ordered levels of desired access to smart-home devices. We chose this list of 24 relationships based on existing users and groups in discretionary access control (DAC) systems and common social relationships in households.

We conducted the pre-study with 31 participants on Amazon’s Mechanical Turk. Participants identified potential concerns for a number of capabilities, in addition to identifying capabilities (e. g., turning on lights) that aroused few concerns. We used these results to generate a list of capabilities, grouping similar functionalities across devices into categories like viewing the current state of a device. We selected the 22 capabilities whose pre-study results showed a spectrum of opinions and concerns while maintaining a feature-set representative of smart homes.

To narrow our initial list of 24 relationships to a tractable number, we examined how pre-study participants assigned each relationship to one of the five ordered categories of desired access to household devices. We chose the six relationships that span the full range of desired access and for which participants were most consistent in their assignments to a category.

5 Methodology

To elicit desired access-control policies for the home IoT, our main study was an online survey-based user study. We recruited participants on Mechanical Turk, limiting the study to workers age 18+ who live in the United States and have an approval rating of at least 95 %.

5.1 Protocol

Each participant was presented with a single capability (e.g., “see which lights in the home are on or off”) randomly chosen from among the 22 identified in the pre-study. Appendix B gives the full list of capabilities and the descriptions participants saw.

We then presented the participant with one of six relationships: spouse; teenage child; child in elementary school; visiting family member; babysitter; neighbor. The text used to describe each relationship is in Appendix C. We first asked whether such a person should be permitted to control that capability “always,” “never,” or “sometimes, depending on specific factors.” These answers were the first step in identifying participants’ desired access-control policies. For the first two options, we required a short free-text justification. To better understand the importance of an authentication method correctly identifying the person in question and the system correctly enforcing the access-control policy, we asked participants who answered “always” or “never” to state how much of an inconvenience it would be if the system incorrectly denied or allowed (respectively) that particular user access to that capability. Participants chose from “not an inconvenience,” “minor inconvenience,” or “major inconvenience,” with a brief free-text justification.

If the participant chose “sometimes,” we required additional explanations to further delineate their desired access-control policy. They first explained in free-text when that person should be allowed to use that capability, followed by when they should not be allowed to do so. On a five-point scale from “not important” to “extremely important,” we asked how important it was for them to have (or not have) access to that capability.

We repeated these questions for the other five relationships in random order. Thus, each participant responded for all six relationships about a single capability.

Afterwards, we asked more general questions about specifying access-control policies for that capability. In particular, we presented eight contextual factors in randomized order, asking whether that factor should influence whether or not anyone should be permitted to use that capability. The possible responses were “yes,” “no,” and “not applicable,” followed by a free-response justification. We asked about the following factors: the time of day; the location of the person relative to the device

(e.g., in the same room); the age of the person; who else is currently at home; the cost of performing that action (e.g., cost of electricity or other monetary costs); the current state of the device; the location of the device in the home; the person’s recent usage of the device. Further, we asked participants to list any additional factors that might affect their decision for that capability.

We concluded with questions about demographics, as well as the characteristics of the participant’s physical house and members of their household. We also asked about their ownership and prior use of Internet-connected devices. Appendix D gives the survey instrument. We compensated participants \$3.50 for the study, which took approximately 20 minutes and was IRB-approved.

5.2 Analysis

Participants’ responses about their access-control preferences included both qualitative free-text responses and multiple-choice responses. Two independent researchers coded the qualitative data. The first researcher performed open coding to develop a code book capturing the main themes, while the second coder independently used that same code book. To quantitatively compare multiple-choice responses across groups, we used the chi-squared test when all cell values were at least 5, and Fisher’s Exact Test (FET) otherwise. For all tests, $\alpha = .05$, and we adjusted for multiple testing within each family of tests using Holm correction.

5.3 Limitations

The ecological validity and generalizability of our study are limited due to our convenience sample on Mechanical Turk. Most of our questions are based on hypothetical situations in which participants imagine the relationships and capabilities we proposed to them and self-report how they expect to react. Furthermore, while some participants were active users of home IoT devices, others were not, making the scenarios fully hypothetical for some participants. We chose to accept this limitation and include recruits regardless of prior experience with home IoT devices to avoid biasing the sample toward early adopters, who tend to be more affluent and tech-savvy.

6 Results

In the following sections we present our findings. We begin by providing an overview of our participants (Section 6.1). Next, we present how desired access-control policies differ across capabilities (RQ1, Section 6.2) and the degree to which desired policies differ across relationships (RQ1, Section 6.3). After that, we show

for which pairs of relationships and capabilities the desired access-control policies are consistent across participants. We use these pairs to derive default policies (RQ2, Section 6.4). Next, we evaluate which contextual factors (e.g., age, location, usage) influence the “sometimes” cases the most, thus explaining users’ reasoning for not always allowing access to a capability (RQ3, Section 6.5). Finally, we analyze the consequences of false authorization and show the impact of falsely allowing / denying access to a certain capability on a per-relationship level (RQ4, Section 6.6).

6.1 Participants

A total of 426 individuals participated in the study, and 425 of them were qualified as effective responses. One response was excluded from our data because their free-text responses were unrelated to our questions. Our sample was nearly gender-balanced; 46 % of participants identified as female, and 54 % as male. The median age range was 25-34 years old (47 %). Most participants (85 %) were between 25 and 54 years old. Some participants (19 %) reported majoring, earning a degree, or holding a job in computer science or a related field.

The majority of our participants (67 %) live in a single-family home, while 25 % live in an apartment. Nearly half of the participants own (49 %) the place where they live, while 47 % rent. Furthermore, we asked how many people (including the participant) live in the same household. Around 20 % of participants reported living in a single-person household, 27 % in a two-person, 23 % in a three-person, and 17 % in a four-person household.

6.2 Capabilities (RQ1)

Current access-control implementation in a smart home system is largely device-based. However, our data motivates a more fine-grained, flexible access-control mechanism. In the following parts, we discuss our main findings, which are visualized in Figure 2.

A) Capability Differences Within a Single Device

We observed that participants’ attitudes toward various capabilities differ within a single device. For example, voice assistants can be used to play music and order things online. However, participants were much more willing to let others play music (32.5 % of participants choose *never* averaged across the six relationships, $\sigma = 0.33$, *median* = 23.7 %) than order things online (59.7 % choose *never* on average, $\sigma = 0.40$, *median* = 71.1 %) (FET, $p < .05$ for the teenager, child, and visiting family member relationships).

Another example of differing opinions across capabilities within a single device include deleting an IoT lock’s activity logs and answering the door, viewing the current

state of the lock, and setting rules for the lock. Across relationships, participants were permissive about capabilities like answering the door (25.6 % chose “never” averaged across all relationships other than children, $\sigma = 0.33$, *median* = 16,7 %). Because children would likely not have a smartphone, we did not ask about them performing this action and we exclude them from this analysis. In contrast, 76.8 % of participants said they would *never* allow others to delete activity logs ($\sigma = 0.28$, *median* = 92.1 %). These differences are significant (FET, all $p < 0.05$ comparing within teenagers, visiting family, and babysitters). Even for a very trust-based relationship like a spouse, some participants still chose *never*. When asked why, one participant wrote: “No one should be able to delete the security logs.”

Even if individuals with relationships like neighbor or babysitter do not live in the same house, permissions are sometimes given when the owner of the house is not around. One typical response for when a capability should be accessible to neighbors is “Perhaps when I’m on vacation and I ask them to watch my home.”

B) Context-Dependent Capabilities

We identified “Answering the Doorbell” to be a highly context-dependent capability. 40 % of participants across relationships ($\sigma = 0.33$, *median* = 38.9 %) selected *sometimes* for this capability. At the same time, an average of 25.6 % of participants across relationships chose *never* ($\sigma = 0.33$, *median* = 16.7 %).

Whether the homeowner is present is a key factor impacting responses. Many participants (66.7 %) chose *sometimes* when it came to the babysitter, because the job itself indicates the parents are not around. If a delivery person rings the doorbell while the babysitter is home, the babysitter should be allowed to handle the event. The majority of participants (77.8 %) also *sometimes* trust a visiting family member with the same level of access. Some participants (16.7 %) will even consider giving this access to their neighbors, so that if there is an emergency when the family is on vacation, their neighbor can see who is at the door from their smartphone.

6.3 Relationships (RQ1)

Relationships play an important role in participants’ preferred access-control policies.

A) Babysitter vs. Visiting Family

In the pre-study, we identified the babysitter and a visiting family member to be members of a guest-like group. In the main study, participants’ overall attitudes toward babysitters and visiting family members were quite consistent with each other. No significant differences are observed between these two relationships in our pairwise chi-squared tests. This is understandable because both




Figure 2: Participants’ desired access-control policies. We introduced participants to a list of relationships (e.g., *neighbor*) and asked them to choose whether someone of that relationship should be permitted to “always,” “sometimes,” or “never” control a capability (e.g., adjust the *camera angle*) in their smart home.

relationships share some trust with the homeowner, while neither lives in the same household.

In general, policies toward a visiting family member are slightly more permissive than policies toward a babysitter. However, analyzing the qualitative data, we found the situation to be more complex. There are some specific capabilities, such as “Live Video,” where babysitters would be granted permissions at a higher rate than a visiting family member. 57.1% of participants decided that a visiting family member would *never* have access to this feature, while only 33.3% of participants decided the same for a babysitter. The reason is that a babysitter’s job is to take care of a child while a parental figure is away. Therefore, the capability itself might help a babysitter take better care of the child, leading to a high rate of granting this permission *sometimes*.

Meanwhile, some features show strong subjective variations, including granting babysitters and visiting family members permission for “Answering the Doorbell.” Some participants found it useful to always allow access, while other participants felt uncomfortable letting someone that is not part of their family have access to this particular capability.

From these observations, we conclude that it is important to have both a relationship-based and capability-based access-control model in a smart home. Such a model should be flexible enough to address the complex needs and use cases that might occur.

B) Child vs. Teenager

Though both children and teenagers are under a parent or guardian’s watch, a teenager (presented as 16 years old) and a child (presented as 8 years old) were given very different access scopes. After removing the five capabilities that are not applicable to a child (whom we assume lacks a smartphone), for twelve of the seventeen remaining capabilities teenagers were given greater access (FET, all $p < .05$). A 16-year-old teenager was regarded as a young adult by many participants and was more widely trusted to use capabilities responsibly. Therefore, the *always* permission was chosen often, and no need for supervision was mentioned in their free-text responses.

Meanwhile, granting an 8-year-old child unencumbered access worried participants much more. Some participants mentioned that they were concerned that a young child would misuse these capabilities, either intentionally or unintentionally, and thus ruin all the settings. Several participants even expressed their worries that a young child could get themselves in danger with the access. For instance, one participant, who selected *never* for the capability of seeing which door is currently locked or unlocked, wrote: “An elementary school child should not be leaving the house on his own accord.” An 8-year-old child’s level of understanding of a smart home system is also questionable. As a result, children rarely were granted access *always* for capabilities other than those related to lights.

Even for capabilities for which participants chose relatively restrictive settings for both teenagers and young children (e.g., “Order Online”), attitudes differed. Though only 5.3% of participants agreed to give full access to “Order Online” to a teenager, 73.7% chose *sometimes* over *never*, giving limited access to their teenager to buy things they needed on Amazon. For young children, 94.7% participants believed that a child at that age should *never* have access to it, frequently justifying that there is no need for younger children to order things online themselves. Many participants mentioned supervision or limitations on what a teenager can buy on Amazon, but they did admit they would let a teenager buy things from Amazon themselves if they had a reason.

C) Overall Preference for Restrictive Policies

We found that, except for spouses and teenagers, most participants preferred a more restrictive access-control policy over a more permissive one. For nine of the twenty-two capabilities averaged over all relationships, more than half of participants chose *never* more frequently than *sometimes*, and *sometimes* more frequently than *always*. Averaged across all capabilities, only 18.1% of participants ($\sigma = 0.12$, *median* = 13.2%) chose *always* for visiting family members, 10.3% for babysitters ($\sigma = 0.09$, *median* = 7.9%), 8.3% for children ($\sigma = 0.10$, *median* = 5.6%) and 0.7% for neighbors ($\sigma = 0.03$, *median* = 0%). There was only a small group of capabilities for which participants were widely permissive: controlling lights and music, which do not have much potential to cause harm or damage.

6.4 Default Policies (RQ2)

In this section, we give an overview of the default deny/allow access policies we observed that capture most participants’ responses. We categorize the policies by relationships and give an in-depth analysis of our findings.

6.4.1 Default Allow

A) Spouses are Highly Trusted

Averaged across all capabilities, 93.5% of participants ($\sigma = 0.09$, *median* = 95.3%) agreed to *always* give access to their spouse, while only 4.15% ($\sigma = 0.05$, *median* = 0%) answered *sometimes*, and 2.35% ($\sigma = 0.06$, *median* = 0%) said *never*. For participants who selected *always*, their most frequent reason was that they fully trust their spouse and that equality should be guaranteed in a marriage. Half of the non-permissive responses came from the capability to delete the smart lock’s log file.

B) Controlling Lights

Access-control policies relating to lights were the most permissive. Looking at the responses for the capability

Table 1: Potential default access-control policies that reflected the vast majority of participants’ preferences.

All
<ul style="list-style-type: none"> • <i>Anyone who is currently at home</i> should always be allowed to adjust lighting • <i>No one</i> should be allowed to delete log files
Spouse
<ul style="list-style-type: none"> • <i>Spouses</i> should always have access to all capabilities, except for deleting log files • <i>No one except a spouse</i> should unconditionally be allowed to access administrative features • <i>No one except a spouse</i> should unconditionally be allowed to make online purchases
Children in elementary school
<ul style="list-style-type: none"> • Elementary-school-age <i>children</i> should <i>never</i> be able to use capabilities <i>without supervision</i>
Visitors (babysitters, neighbors, and visiting family)
<ul style="list-style-type: none"> • <i>Visitors</i> should only be able to use any capabilities <i>while in the house</i> • <i>Visitors</i> should <i>never</i> be allowed to use capabilities of <i>locks, doors, and cameras</i> • <i>Babysitters</i> should only be able to <i>adjust the lighting and temperature</i>

to turn lights on and off, most responses align with a proposed default policy of people only being able to control the lights if they are physically present within the home. Relatedly, some participants chose *sometimes* for visiting family members and babysitters, depending on whether they are physically present within the home.

6.4.2 Default Deny

A) Lock Log Sensitivity

As mentioned in Section 6.2, “Delete Lock Log” is the capability least frequently permitted, and access should therefore be denied by default. Even for a spouse, this capability should not be accessed by default (only 68.4% chose *always* for their spouse). More than 75% of participants chose *never* for all other relationships. As the main method of retrospectively using history, the log is not meant to be deleted.

B) Supervising Children

The elementary-school-age child (presented as 8 years old) was one of the most restricted relationships. On average across all capabilities, 69.4% of participants chose *never* for the child ($\sigma = 0.19$, *median* = 70.6%). Only neighbors received fewer permissions. In our chi-squared tests, we did not observe significant differences in desired access-control settings for children between participants who are currently living with a child, who have lived with a child before, and who have never lived with a child. None of our capabilities were considered child-friendly enough for even the majority of participants to *always* grant their elementary-school-age child

access to that capability *always*. For only the “Light State” and “Play Music” capabilities was *never* chosen by fewer than half of participants. Despite being an immediate family member and living together, plenty of participants expressed fears that a child at that age might toy with these features and unintentionally mess up their settings or even cause danger to themselves. With supervision, though, many participants would consider giving temporary access to their children to gradually teach them how to use such a new technology.

C) Ordering Online

The capability to make an online purchase was generally limited to spouses only; 78.9% of participants said that only their spouse should always be allowed to make online purchases, but 84.2% also said that it was acceptable for non-spouse users to do the same if given explicit permission by the homeowner.

D) Administrative Capabilities

By default, only spouses should be able to access administrative capabilities, such as adding users, connecting new devices, and installing software updates. 89.7% of participants gave their spouse access to these administrative capabilities *always*, while only 39.7% of participants *always* gave comparable access to their teenage child. Unsurprisingly, under twenty percent of participants would give full access to other relationships.

6.5 The Impact of Context (RQ3)

Since there are many factors at play in the access-control-policy specification process, it is important to identify which contextual factors are most influential in this process and how they contribute to the final decision. The full results are visualized in Figure 3. We also ran chi-squared tests to see if each contextual factor had a relatively greater influence on some capabilities rather than others. While we did not observe significant differences for the “People Nearby”, “Cost” and “Usage History” contextual factors across capabilities, we did observe significant differences for the other five contextual factors.

A) Age

The *age* of the user was the most influential factor on average across the eight capabilities (78.1% on average, $\sigma = 0.13$, *median* = 78.3%), and the proportion of participants for whom age mattered varied across capabilities ($p = 0.040$). The main capability for which age played less of a role was for *changing the camera angle* (only 50%). Many participants were concerned with letting a young person have access to certain capabilities. “*They need to be mature enough to use it responsibly*” was one typical response. However, another participant instead explained, “*It will be the person themselves and how capable they are with technology. I do not care about age.*”. Thus, while *age* was frequently mentioned,

in reality the decision process is more likely to be driven by how capable and responsible a user is, which sometimes correlates with the user’s age. Our results indicate that a child at a young age (around 8 years old) is generally not perceived to be tech-savvy and responsible enough to be allowed unsupervised access.

B) Location of Device

The proportion of participants for whom the device’s location impacted the access-control policy varied across capabilities ($p < 0.001$). Capabilities relating to cameras were unsurprisingly very location-sensitive. “Camera Angle” is the only capability for which a device’s *location* was more frequently influential (70% of participants) than the user’s *age*. *Device location* was the second most frequently invoked factor for turning a camera on or off (60%) and watching live video (81%). If a smart camera is installed indoors, especially in a bedroom or bathroom, it will be much more privacy-sensitive. Participants reflected this by saying, for example, “*I can see where a guest/house-sitter/baby-sitter might need to access a view of outside or the garage but not inside.*” Therefore, when designing a smart camera, whether the camera will be used indoors or outdoors should be considered and reflected in default access-control policies.

C) Recent Usage History

The proportion of participants for whom a device’s recent usage history impacted their access-control policy did not differ significantly across capabilities. On average across capabilities, 51.7% of participants ($\sigma = 0.12$, *median* = 52.6%) agreed that this factor impacted their decision about the access-control policy. For participants who felt the device’s recent usage history would change their decision, two main rationales arose. On the one hand, if the history states that a user is abusing a capability, then the owner may revoke access. One participant wrote, “*If someone were to misuse the device, you best bet they aren’t getting a second chance. Alright maybe I’ll give them a second chance, but definitely not a third!*”. On the other hand, if a user turns out to be trustworthy, then the owner may consider letting them keep the access, or even extending it. “*If my kid had been using the device responsibly, I would feel more comfortable giving them more access.*” However, some participants felt the recent usage history was not particularly relevant for two main reasons. First, if the involved capability itself cannot cause much trouble, such as “Light Scheme,” a common line of reasoning is that “*It would be hard to abuse this capability, so it doesn’t matter to me.*” Second, if the capability itself is so concerning that participants are reluctant to give others access (e.g., “Delete Video”), usage history did not play a role.

D) Time of Day

The importance of the *time of day* contextual factor

Impact of Contextual Factors on Capabilities

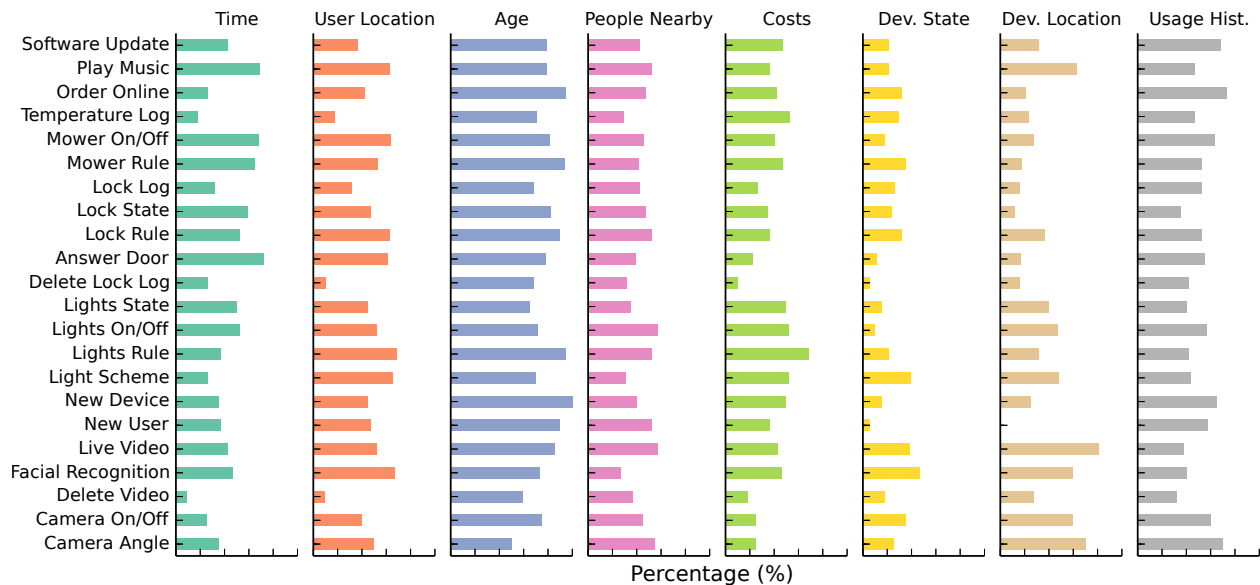


Figure 3: Contextual factors: Sometimes access must depend on the context. In the study we asked participants for such factors and identified multiple that are very influential (such as the age of the user) and learned how they contribute to the decision make process.

varied across capabilities ($p = 0.001$). “Play music” (68.4 %) and lawnmower-related capabilities (64.7 % for creating rules for the mower, 68.2 % for turning lawn mower on/off remotely) were particularly sensitive to the time of the day. In order to not interrupt other people’s rest, participants tended to limit lawnmower usage to the daytime and playing music to the early evening.

E) Location of User

Capabilities that change devices’ behaviors tended to be more sensitive to where the user is physically located when trying to control the device ($p < 0.001$) since many functionalities cannot be enjoyed without proximity. For example, creating rules that control the lights (68.4 % of participants felt the user’s location mattered) and “Facial Recognition” (66.7 %) were prime examples. Many participants wrote that they would not want anyone who is not currently present in the house to use these capabilities unless it is the owner or their spouse.

F) Costs

The influence of the cost of exercising a capability did not vary across capabilities ($p = 0.162$). We believe this is in part due to our study design that did not include high-wattage appliances. Nevertheless, we observed some evidence of concerns with the cost of leaving lower-wattage devices, like lights, on during the day. Some participants mentioned that while lights do not consume a lot of electricity, cost can quickly become a concern if heavy appliances were to be involved. In ad-

dition, the influence of cost on online shopping differed due to different interpretations of cost. For cases where participants did indicate that cost is a concern, their interpretation was based on the cost of the good purchased, rather than the electricity used in placing an order.

G) People Nearby

43.6% of participants ($\sigma = 0.09$, $median = 43.6\%$) indicated that who else is nearby might impact their access-control decision. The role of people nearby did not differ significantly across capabilities ($p = 0.400$). For participants who believe this factor matters, there are two contrasting conclusions. Some people might feel more permissive when they themselves are around since that means they can supervise everything. However, others felt less permissive because if they are around, there is no need for others to have access since the others simply would need to ask the owner. Therefore, it is important for the system configuration to take these divergent mental models into consideration, letting users decide which direction they might choose to go in.

H) State of Device

The current state of device was overall the least important factor in participants’ access-control decisions on average ($mean = 23.7\%$, $\sigma = 0.11$, $median = 22.3\%$), though this importance did differ across capabilities ($p = 0.044$). Notably, 46.7 % of participants who answered about the “Facial Recognition” the capability marked the state of the device as an influential factor. This is because

if the camera is currently off, then there is no reason for anyone to enable or disable the facial recognition.

I) Other Factors

We included a free-text question with which participants could list other factors they thought played a role in their access-control-policy specification process. In their responses, we observed a long tail of additional contextual factors, including weather, people's familiarity with technology, how close they are to the owners, and the frequency of one's access to a certain capability.

6.6 Wrong Decisions' Consequences (RQ4)

Analyzing consequences of incorrect authorization decisions, we can learn how much tolerance a user has for a policy to fail given a specific capability and relationship pair. It is crucial to understand how strongly users would feel if the system were to malfunction. We analyze *false allow* and *false deny* decisions separately.

6.6.1 False Allow

Note that responses about *falsely allowing access* belong to those participants who intended never to grant access to a certain capability to a certain relationship. These participants therefore might be more concerned than other participants in certain aspects, which leads to some narrow tensions with the broader trends seen in previous sections. Figure 4 (top) summarizes these results.

A) Neighbor false allows a major inconvenience

Across all capabilities, 64.1 % of the participants stated that it is a *major inconvenience* if the authorization system gives access to their neighbor by accident. Turning the security camera on or off (100 % a major inconvenience) and creating rules for a smart lock (92.9 % a major inconvenience and 7.1 % a minor inconvenience) are the most concerning capabilities. Note that in the study, we described the people representing the relationship *neighbor* as "good people, which includes friendly small talk, and occasional dinner invitations." Nevertheless, privacy and security were major concerns.

B) Spousal false allows have severe consequences

Though the number of false-allow responses for the spouse relationship is quite small ($n = 10$), it still gives some interesting insights. 50 % of the answers are based on deleting log files from a smart lock. Four out of five respondents rate falsely allowing a spouse to delete the log file not to be an inconvenience. "*I wouldn't really care about my spouse deleting it, but it would bother me that the system is not secure,*" was a typical response. There were five more responses from other capabilities. From those, four out of five indicated that a false allow decision was a major inconvenience. It is surprising to see that a few participants believed it a major issue if the

mechanism allows their spouse to access certain capabilities by mistake.

C) Visiting family false allows a minor issue

Though we presented earlier that participants' permissiveness toward a visiting family member and a babysitter was very similar (and tended toward not being permissive), we observed a distinction when it comes to false allows. Participants were much less concerned with incorrectly giving access to a visiting family member (70 % chose *minor* or *not an inconvenience*) than to a babysitter (58 %). Responses like "*He is my family member so I trust him a bit*" were common. While participants believed the visiting family member would not do much harm, false allows would still upset them a bit.

D) Shopping / lawn mowers forbidden for children

Among all capabilities, incorrectly allowing a young child to order online (79 % a *major inconvenience*) and create rules for the lawn mower (70.6 %) were the two capabilities where false allows for a child raised great concern. A child at such a young age is generally not trusted with ordering things online. "*The child could spend a ton of money on products we don't need,*" wrote one participant. A lawn mower is considered dangerous. One participant simply wrote, "*(A lawn mower) could cause harm to the child.*"

6.6.2 False Deny

Responses in this section, *falsely denying access*, come from participants who intended to give access to a certain relationship. Figure 4 (bottom) visualizes the full results.

A) Participants Did Not Want to be Locked Out

Lock-related capabilities raised the most concern (63.9 % of responses for "Lock State" and 58.8 % for "Lock Rule" found falsely denying access *major inconveniences*). Participants tended to be very cautious about smart locks. Even though viewing a lock state does not directly concern locking or unlocking the door, participants still worried whether a malfunctioning access-control system would lock people out, thus marking these false denials as major inconveniences.

B) Spouses and Trust Issues

One common reason why participants gave full access to their spouse is because they believe two people in a marriage should be equal, which means two parties should have the same access to a system. Therefore, if their spouse is accidentally rejected by the system, it could raise trust issues and spur arguments within the marriage. We found a number of responses similar to "*I would not want my spouse to think I did not trust them.*" It is interesting to see that not only do relationships impact access-control policies, but relationships are also influenced by authorization results. Thus, extra care is required for such relationships.

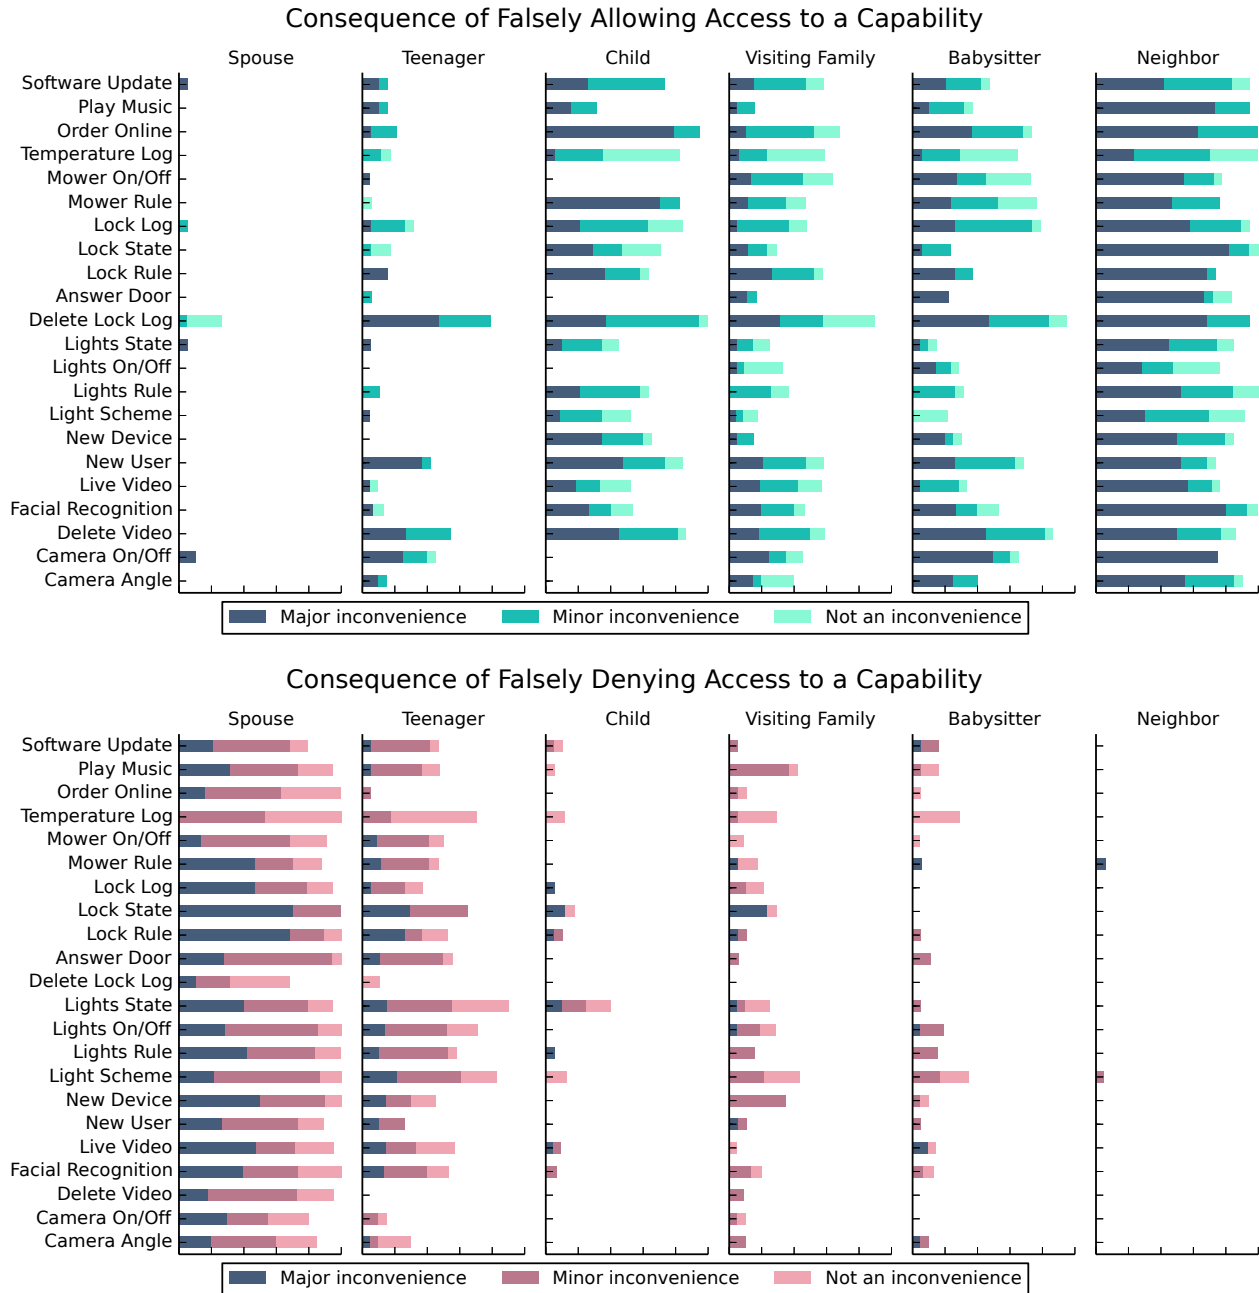


Figure 4: Perceived consequences of incorrectly allowing someone to use a capability when they should never be permitted to do so (top) or incorrectly denying someone when they should always be permitted to do so (bottom).

7 Discussion

Capabilities, Relationships, and Context. While access control in smart homes is currently often device-centric, our user study demonstrated that a capability- and relationship-centric model more closely fits user expectations. Home IoT technologies allow for multiple ways of achieving the same end result, whereas devices often bring together vastly different capabilities. For ex-

ample, to increase a room’s brightness, one could remotely turn on a light using a smartphone app, remotely open the shades, or ask a voice assistant to do either. This model reveals nuances that are missed in the device-centric model. From the data for RQ1, we see that the desired policies can vary widely within a *single* device based on the relationship and the context of access. Although some of these distinctions are intuitive (e.g., child vs. teenager), others are more nuanced and surpris-

ing (e. g., babysitter vs. visiting family member). They also provide a concrete access-control vocabulary for developers of future smart-home devices.

A difficult decision in access-control systems involves default policies. In multi-user social environments, intuition suggests a default policy would be complex. Surprisingly, our data for RQ2 suggests that potential default policies are actually simple and reminiscent of non-IoT policies. For example, our default policy says that a person can actuate a light if they are physically close to it. Though IoT lights can be remotely actuated, the relation between proximity and using a light is not broken. Although conceptually simple, this rule's enforcement is non-trivial, requiring creating and deploying authentication methods beyond the possession of a smartphone.

Data from RQ3 suggests that the factors affecting access-control decisions are heavily context-dependent. Current home IoT devices only support rudimentary forms of context (Section 2). Some contextual factors, such as age, are currently present in smartphones and cloud services (e. g., *Apple's iCloud Family Sharing* supports adding a child Apple ID that requires parental approval for purchases, while Netflix has *kids* option). We recommend that for home IoT settings, these contextual factors should be a first-order primitive.

Based on these findings (RQ1-3), we envision several changes to smart-home setup. This process currently involves installing hubs and devices with a set of coarse-grained accounts. Our work suggests that future smart homes could instead set access-control policies by walking users through a questionnaire whose vocabulary derives from our user study. This is closer to the experience of setting up software, where a package comes with secure defaults that are customized to the specific installation. Using default policies derived from our results would minimize user burden since it would reflect common opinions by default. Physical control (e.g., switches) already enables certain default policies, so software authorization might seem unnecessary in certain situations. However, switches are often add-ons to IoT starter kits, making software authorization a prerequisite to a satisfying user experience.

Authorization Vocabulary. Based on our study results, we discuss a potential authorization vocabulary that is helpful in building future authorization and authentication for home IoT platforms. The basic unit of the vocabulary is a triplet containing \langle Capability, UserType, Context \rangle . As discussed, capabilities better capture the nuances of access control in the home than devices. Appendix A lists capabilities commonly supported by current home IoT platforms. UserType captures the relationship of the user to the home, and to the owners. From our study, these types should nonexhaustively include: Spouse, Teenager, Child, Babysitter, and Neighbor.

Spouses tend to be users with the highest levels of access, generally equivalent to administrators in traditional computing systems. Context refers to the environmental factors that might affect an access-control decision. For example, certain parents might be more permissive in allowing a child to watch TV without supervision. Based on our study, at the minimum context should include: Time, User Location, Age, People Nearby, Cost of Resource, Device State, Device Location, and Usage History. Depending on the Capability and the UserType components of the triplet, the importance of the context can change. For example, for a UserType of Child, the 'People Nearby' contextual factor plays a prominent role in the access-control decision. However, for spouses, it generally has no bearing. The same goes for the Capability. The 'Device Location' contextual factor is crucial for camera-related capabilities, but not so important for the capability of adding a new user.

Mapping Authorization and Authentication. Although we focused on analyzing access control, we briefly discuss how our findings affect the design of authentication mechanisms. Below, we discuss a set of authentication mechanisms and comment on their ability to identify users, relationships, and contextual factors. We also discuss privacy limitations and the effect of false positive and negatives.

Smartphones are the most widely used devices to access IoT devices in the home. Users may present their identity to a device using a password, PIN, or (more recently) fingerprints. These identities can be used by home IoT devices to determine the identity, and hence relationship, of the person attempting access. From the perspective of false positives/negatives, smartphones can closely match user expectations. They are inconvenient, however, for temporary visitors because they require the visitor to install an app and the owner to authorize them.

Wearable devices like watches, glasses, and even clothing [18] might serve as proxy devices with more natural interactions than a smartphone. For example, a user can gesture at a nearby device to control it (e. g., wave at a light to turn it on or off). As each user will perform a gesture differently, it can also serve as a form of authentication and thus be used to identify a person and their relationship. Furthermore, the proximity of a wearable device is helpful in identifying several contextual factors, including user location and nearby people. From a false positive/negative perspective, biometrics require quite a bit of tuning that can affect an owner's choice of using this method, especially when authenticating high-access spouses or for operating dangerous equipment like lawn mowers.

Voice assistants are increasingly ubiquitous in homes. Although such assistants can perform speaker identification (e.g., Google Home Voice Match), they are cur-

rently used as a personalization hint rather than a security boundary. However, future versions that use additional hardware might be useful in determining a speaker's identity and relationship for access-control purposes [12]. Such assistants could help identify contextual factors like the location of a user or the presence of nearby people (e.g., a supervising adult near children). From the perspective of false positives/negatives, any voice-based method will require tuning. Audio is especially sensitive to background noise. Audio authentication also introduces privacy issues, as well as the potential for eavesdropping and replay attacks.

Advances in computer vision can also be leveraged to identify users, their relationship, and their location within a home with cameras. However, it is possible for computer vision systems to falsely identify individuals or confuse identities. Thus, some level of false positive/negative tuning will be required, especially when a household is expected to have many temporary occupants. A big downside of this mechanism is the privacy risk—cameras can track home activity at a high level of granularity. However, some of the privacy issues could potentially be alleviated using local processing or privacy-preserving vision algorithms [21].

Bilateral or continuous authentication mechanisms embody the idea that a user has to be: (a) physically present, and (b) currently using the device [20, 28]. Such mechanisms are readily able to identify users and relationships, and to support contextual factors involving user presence. False positive/negative tuning varies based on the specific instantiation. If a wearable device with a continuous authentication algorithm is used, then the false positive/negative rates must be considered. Privacy concerns can be alleviated if this mechanism is implemented in a decentralized manner—only the user's proxy device and the target device are involved in establishing an authenticated channel. It can also provide a simple solution to the de-authentication problem (revoking access if a temporary visitor is no longer welcome).

In sum, we have taken initial steps toward reenvisioning access-control specification and authentication in the home IoT. Much work remains in continuing to translate these observations to fully usable prototypes, as well as in supporting ever richer capabilities and interactions.

8 Acknowledgments

We thank the reviewers and our shepherd, Adam Bates, for their insightful feedback, as well as Camila Cuesta Arcentales for help on the study instrument. This material is based upon work supported by the National Science Foundation under Grants No. 1756011 and 1565252. Earlene Fernandes was supported by the UW Tech Policy Lab and the MacArthur Foundation. Maximilian

Golla was supported by the German Research Foundation (DFG) Research Training Group GRK 1817/1.

References

- [1] AL-MUHTADI, J., RANGANATHAN, A., CAMPBELL, R., AND MICKUNAS, M. D. A Flexible, Privacy-Preserving Authentication Framework for Ubiquitous Computing Environments. In *Proc. ICDCS* (2002).
- [2] AMAZON. Echo, Nov. 2014. <https://www.amazon.com/echo>, as of June 29, 2018.
- [3] ANTONAKAKIS, M., APRIL, T., BAILEY, M., BERNHARD, M., BURSZTEIN, E., COCHRAN, J., DURUMERIC, Z., HALDERMAN, J. A., INVERNIZZI, L., KALITSIS, M., KUMAR, D., LEVER, C., MA, Z., MASON, J., MENSCHER, D., SEAMAN, C., SULLIVAN, N., THOMAS, K., AND ZHOU, Y. Understanding the Mirai Botnet. In *Proc. USENIX Security Symposium* (2017).
- [4] BAUER, L., CRANOR, L. F., REEDER, R. W., REITER, M. K., AND VANIEA, K. Real Life Challenges in Access-control Management. In *Proc. CHI* (2009).
- [5] BELKIN. WeMo Home Automation, Jan. 2012. <https://www.belkin.com/wemo>, as of June 29, 2018.
- [6] BONNEAU, J., HERLEY, C., VAN OORSCHOT, P. C., AND STAJANO, F. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *Proc. IEEE SP* (2012).
- [7] BRUSH, A. B., JUNG, J., MAHAJAN, R., AND MARTINEZ, F. Digital Neighborhood Watch: Investigating the Sharing of Camera Data Amongst Neighbors. In *Proc. CSCW* (2013).
- [8] DENNING, T., KOHNO, T., AND LEVY, H. M. Computer Security and the Modern Home. *CACM* 56, 1 (2013), 94–103.
- [9] ENCK, W., ONGTANG, M., AND MCDANIEL, P. Understanding Android Security. *IEEE Security & Privacy* 7, 1 (2009), 50–57.
- [10] FELT, A. P., EGELMAN, S., AND WAGNER, D. I've Got 99 Problems, but Vibration Ain't One: A Survey of Smartphone Users' Concerns. In *Proc. SPSM* (2012).
- [11] FELT, A. P., HA, E., EGELMAN, S., HANEY, A., CHIN, E., AND WAGNER, D. Android Permissions: User Attention, Comprehension, and Behavior. In *Proc. SOUPS* (2012).
- [12] FENG, H., FAWAZ, K., AND SHIN, K. G. Continuous Authentication for Voice Assistants. In *Proc. MobiCom* (2017).
- [13] FERNANDES, E., JUNG, J., AND PRAKASH, A. Security Analysis of Emerging Smart Home Applications. In *Proc. IEEE SP* (2016).
- [14] FERNANDES, E., PAUPORE, J., RAHMATI, A., SIMONATO, D., CONTI, M., AND PRAKASH, A. FlowFence: Practical Data Protection for Emerging IoT Application Frameworks. In *Proc. USENIX Security Symposium* (2016).
- [15] FERNANDES, E., RAHMATI, A., EYKHOLT, K., AND PRAKASH, A. Internet of Things Security Research: A Rehash of Old Ideas or New Intellectual Challenges? *IEEE Security & Privacy* 15, 4 (2017), 79–84.
- [16] FOULADI, B., AND GHANOUN, S. Honey, I'm Home!!, Hacking ZWave Home Automation Systems, July 2013. Black Hat USA.
- [17] GOOGLE. Android Supporting Multiple Users, June 2017. <https://source.android.com/devices/tech/admin/multi-user>, as of June 29, 2018.
- [18] GOOGLE. Jacquard Powered Smart Jackets, Sept. 2017. <https://atap.google.com/jacquard/>, as of June 29, 2018.

- [19] GOOGLE. Set up Voice Match on Google Home, Oct. 2017. <https://support.google.com/googlehome/answer/7323910>, as of June 29, 2018.
- [20] HUHTA, O., UDAR, S., JUUTI, M., SHRESTHA, P., SAXENA, N., AND ASOKAN, N. Pitfalls in Designing Zero-Effort Deauthentication: Opportunistic Human Observation Attacks. In *Proc. NDSS* (2016).
- [21] JANA, S., NARAYANAN, A., AND SHMATIKOV, V. A Scanner Darkly: Protecting User Privacy from Perceptual Applications. In *Proc. IEEE SP* (2013).
- [22] JIA, Y. J., CHEN, Q. A., WANG, S., RAHMATI, A., FERNANDES, E., MAO, Z. M., AND PRAKASH, A. ContextIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *Proc. NDSS* (2017).
- [23] JOHNSON, M., AND STAJANO, F. Usability of Security Management: Defining the Permissions of Guests. In *Proc. SPW* (2006).
- [24] KIM, T. H.-J., BAUER, L., NEWSOME, J., PERRIG, A., AND WALKER, J. Challenges in Access Right Assignment for Secure Home Networks. In *Proc. HotSec* (2010).
- [25] KIM, T. H.-J., BAUER, L., NEWSOME, J., PERRIG, A., AND WALKER, J. Access Right Assignment Mechanisms for Secure Home Networks. *Journal of Communications and Networks* 13, 2 (2011), 175–186.
- [26] LEKAKIS, V., BASAGALAR, Y., AND KELEHER, P. Don’t Trust Your Roommate or Access Control and Replication Protocols in “Home” Environments. In *Proc. HotStorage* (2012).
- [27] LIU, J., XIAO, Y., AND CHEN, C. P. Authentication and Access Control in the Internet of Things. In *Proc. ICDCS* (2012).
- [28] MARE, S., MOLINA-MARKHAM, A., CORNELIUS, C., PETERSON, R., AND KOTZ, D. ZEBRA: Zero-Effort Bilateral Recurring Authentication. In *Proc. IEEE SP* (2014).
- [29] MATTHEWS, T., O’LEARY, K., TURNER, A., SLEEPER, M., WOELFER, J. P., SHELTON, M., MANTHORNE, C., CHURCHILL, E. F., AND CONSOLVO, S. Stories from Survivors: Privacy & Security Practices when Coping with Intimate Partner Abuse. In *Proc. CHI* (2017).
- [30] MAZUREK, M. L., ARSENAULT, J. P., BRESEE, J., GUPTA, N., ION, I., JOHNS, C., LEE, D., LIANG, Y., OLSEN, J., SALMON, B., SHAY, R., VANIEA, K., BAUER, L., CRANOR, L. F., GANGER, G. R., AND REITER, M. K. Access Control for Home Data Sharing: Attitudes, Needs and Practices. In *Proc. CHI* (2010).
- [31] NAEINI, P. E., BHAGAVATULA, S., HABIB, H., DEGELING, M., BAUER, L., CRANOR, L. F., AND SADEH, N. Privacy Expectations and Preferences in an IoT World. In *Proc. SOUPS* (2017).
- [32] PHILIPS. Hue, Oct. 2012. <https://www.meethue.com>, as of June 29, 2018.
- [33] PROSPERO, M. Best Smart Home Gadgets of 2018, Jan. 2018. <https://www.tomsguide.com/us/best-smart-home-gadgets,review-2008.html>, as of June 29, 2018.
- [34] PULLEN, J. P. Amazon Echo Owners Were Pranked by South Park and Their Alexas Will Make Them Laugh for Weeks, Sept. 2017. <http://fortune.com/2017/09/14/watch-south-park-alexa-echo/>, as of June 29, 2018.
- [35] SAMSUNG. SmartThings: Add a Little Smartness to Your Things, Aug. 2014. <https://www.smartthings.com>, as of June 29, 2018.
- [36] SAMSUNG. SmartThings: Capabilities Reference, Jan. 2018. <https://smartthings.developer.samsung.com/develop/api-ref/capabilities.html>, as of June 29, 2018.
- [37] SCHAUB, F., BALEBAKO, R., DURITY, A. L., AND CRANOR, L. F. A Design Space for Effective Privacy Notices. In *Proc. SOUPS* (2015).
- [38] SCHECHTER, S. The User IS the Enemy, and (S)he Keeps Reaching for that Bright Shiny Power Button! In *Proc. HUPS* (2013).
- [39] STOBERT, E., AND BIDDLE, R. Authentication in the Home. In *Proc. HUPS* (2013).
- [40] SURBATOVICH, M., ALJUR AidAN, J., BAUER, L., DAS, A., AND JIA, L. Some Recipes Can Do More Than Spoil Your Appetite: Analyzing the Security and Privacy Risks of IFTTT Recipes. In *Proc. WWW* (2017).
- [41] TIAN, Y., ZHANG, N., LIN, Y.-H., WANG, X., UR, B., GUO, X., AND TAGUE, P. SmartAuth: User-Centered Authorization for the Internet of Things. In *Proc. USENIX Security Symposium* (2017).
- [42] TILLEY, A. How A Few Words To Apple’s Siri Unlocked A Man’s Front Door, Sept. 2016. <https://www.forbes.com/sites/aarontilley/2016/09/21/apple-homekit-siri-security>, as of June 29, 2018.
- [43] UR, B., JUNG, J., AND SCHECHTER, S. The Current State of Access Control for Smart Devices in Homes. In *Proc. HUPS* (2013).
- [44] UR, B., JUNG, J., AND SCHECHTER, S. Intruders Versus Intrusiveness: Teens’ and Parents’ Perspectives on Home-entryway Surveillance. In *Proc. UbiComp* (2014).
- [45] WANG, Q., HASSAN, W. U., BATES, A., AND GUNTER, C. Fear and Logging in the Internet of Things. In *Proc. NDSS* (2018).
- [46] WIJESEKERA, P., BAO KAR, A., HOSSEINI, A., EGELMAN, S., WAGNER, D., AND BEZNOSOV, K. Android Permissions Remystified: A Field Study on Contextual Integrity. In *Proc. USENIX Security Symposium* (2015).
- [47] WONG, V. Burger King’s New Ad Will Hijack Your Google Home, Apr. 2017. <https://www.cnn.com/2017/04/12/burger-kings-new-ad-will-hijack-your-google-home.html>, as of June 29, 2018.
- [48] YANG, R., AND NEWMAN, M. W. Learning from a Learning Thermostat: Lessons for Intelligent Systems for the Home. In *Proc. UbiComp* (2013).
- [49] YU, T., SEKAR, V., SESHAN, S., AGARWAL, Y., AND XU, C. Handling a Trillion (Unfixable) Flaws on a Billion Devices: Rethinking Network Security for the Internet-of-Things. In *Proc. HotNets* (2015).
- [50] ZENG, E., MARE, S., AND ROESNER, F. End User Security and Privacy Concerns with Smart Homes. In *Proc. SOUPS* (2017).
- [51] ZHANG, N., DEMETRIOU, S., MI, X., DIAO, W., YUAN, K., ZONG, P., QIAN, F., WANG, X., CHEN, K., TIAN, Y., GUNTER, C. A., ZHANG, K., TAGUE, P., AND LIN, Y. Understanding IoT Security Through the Data Crystal Ball: Where We Are Now and Where We Are Going to Be. *CoRR abs/1703.09809* (2017).

APPENDIX

A Home IoT Devices Considered

Cooking Devices	Anova Culinary Precision Cooker Char-Broil Digital Electric Smoker June Intelligent Oven Perfect Bake Pro Samsung Family Refrigerator
Hubs	Samsung SmartThings Wink Hub 2
Lights/Power Plugs	Belkin Wemo Insight Switch BeOn iHome Smartplug LIFX Color 1000 Lutron Caseta In-Wall Wireless Lighting Philips Hue Starter Set
Locks	August SmartLock Kwikset Smartcode Touchscreen
Outdoor Devices	Rachio Smart Sprinkler Robomow
Security Cameras	Kuna Toucan LG Smart Security Wireless Camera Nest Cam NetGear ArloPro Skybell Video Doorbell Tend Secure Lynx Indoor
Thermostats	EcoBee 4 Hisense Portable AC Nest Learning Thermostat
Voice Assistants	Amazon Echo Echo Dot Google Home

B Full Descriptions of Capabilities

- **Software Update:** Install a software update to get the latest features, improvements, and security updates.
- **Play Music:** Play music (e. g., from Spotify) in the house.
- **Order Online:** Make online purchases (e. g., on Amazon) on a shared household account.
- **Temperature Log:** View the last 10 temperature adjustments and who made them.
- **Mower On/Off:** Turn the lawn mower on or off remotely (i. e., on a smartphone, from anywhere).
- **Mower Rule:** Create rules that specify what the lawn mower should do, connecting its actions to other devices, sensors, and services. For example, one could create a rule specifying that the mower should not mow if it is raining.
- **Lock Log:** View an activity log for the past week that shows who entered the home at what times. People will be identified based on whose PIN code or smartphone was used to unlock the door.
- **Lock State:** See whether the front door is currently locked or unlocked.
- **Lock Rule:** Create rules that specify when the lock should be locked or unlocked, connecting it to other devices, sensors, and services. For example, one could create a rule specifying that the lock should always be locked when no one is home.
- **Answer Door:** Answer the doorbell by seeing a live video of who is at the front door and having the opportunity to unlock the door remotely (e. g., on a smartphone, from anywhere).
- **Delete Lock Log:** Delete the activity log that records who has tried to open or close the door.

- **Lights State:** See which lights in the home are on or off.
- **Lights On/Off:** Remotely control whether a light is currently on, as well as how bright it is (e. g., on a smartphone, from anywhere).
- **Lights Rule:** Create rules that specify when the lights should turn on/off or change color based on other sensors, devices, and services. For example, one could create rules specifying how the lights automatically change brightness or color based on the current weather or the movie played on the TV.
- **Light Scheme:** Allow a streaming video provider to change the lighting according to the theme of the movie that is currently being watched.
- **New Device:** Connect a new device to the hub, enabling the hub to control that device.
- **New User:** Add new users (people) to the smart-home management system, as well as remove users from the smart-home management system.
- **Live Video:** See live video from each camera in or around the house.
- **Facial Recognition:** Enable or disable facial recognition technology for a person. This technology is used to identify them automatically in video recordings.
- **Delete Video:** Delete one or more previously recorded videos.
- **Camera On/Off:** Turn the camera on/off remotely (e. g., on a smartphone, from anywhere).
- **Camera Angle:** Change camera's view remotely (including turning its lens to view a different angle, zooming in/out, etc.).

C Full Descriptions of Relationships

- **Your spouse:** Imagine you have a spouse. You live with them everyday and share all smart appliances in your home. You make decisions together in most cases, especially important ones.
- **Your teenage child:** Imagine you have a 16-year-old child. They live with you, go to school in the morning, and come back in the afternoon (on the weekdays). They are familiar with all of these Smart devices in your home, and enjoy using them. They know how to use these devices as well as you do, if not better. They spend a lot of time on their smartphone. They usually are well-behaved, but they are still a teenager.
- **Your child in elementary school:** Imagine you have an 8-year-old child who is still in elementary school. They live with you and go to school daily, unless it's the weekend or a holiday. They have a basic idea of how to use smart devices. However, they don't know how to use some more complex features properly, like changing the settings, but it doesn't discourage them from trying. They do not have their own smartphone, but they keep asking you for one.
- **A visiting family member:** Imagine you have a visiting family member. They are about the same age as you, if not much older. You grew up together, but now you meet each other once or twice a year, because you live far away from each other. They visit you on holidays or other big events. They usually stay with you for several days, maybe even a little bit past the holiday, and they remain at home alone while you are away for work.
- **The babysitter:** Imagine you have a babysitter in your home for taking care of your child. They will be at your place while you are at work. They work 4 hours after school, 3 days per week. You have known them over 6 months and you are satisfied with their work so far, and have no intention of letting them go anytime soon.
- **Your neighbor:** Imagine you have a neighbor living next to you. You don't know them very well, but they seem to be good people. If you meet them on the street, you greet them and make some friendly small talk. Occasionally you invite them over for dinner, but they are never in your house when you are away.

D Survey Instrument

Introduction Computing is transitioning from single-user devices, such as laptops and phones, to the Internet of Things, in which many users will interact with a particular device, such as an Amazon Echo or Internet-connected door lock. Current measures fail to provide usable authentication, access control, or privacy when multiple users share a device. Even more so, the users of a given device often have complex social relationships to each other. Our goal is to develop techniques and interfaces that enable accurate access control and authentication in multi-user IoT environments, based on user preferences.

Participation should take about 20 minutes.

In recent years, many internet-connected (“smart”) home devices and appliances have entered the market. Imagine that you own many such smart devices that are connected both to the Internet and to each other.

This includes a smart hub that can control other devices in your home, particularly with the help of the smart voice assistant. You also have a smart door lock and smart camera for home security, as well as smart lighting and a smart thermostat to control your environment. There is also a smart lawn mower maintaining your lawn. All of these devices can be remotely controlled using a smartphone app by anyone to whom you have given permission. You, or anyone else you have permitted, can also write rules specifying in what situations devices should activate automatically.

In this survey, we will ask you questions about who in your household should be allowed to access one particular feature of a smart device. If you live in multiple places, think of the home in which you live the majority of the time. For all questions, assume that the system has correctly identified the user involved (i. e., there are no cases of mistaken identity).

Because the situations may involve either positive or negative consequences, you should take some time to think about your response. The next button will not appear until you have spent at least 30 seconds on each page.

In this survey, we will ask whether you will allow people of the following relationships to control a particular feature of a smart device: your spouse; your teenage child; your child in elementary school; a visiting family member; a babysitter; your neighbor. Please imagine you have these relationships in your life even if you don’t. All of these relationships are separate people.

If you grant access to any of these people, they will be able to access your devices whether or not they are in your home, unless you specify otherwise in your responses in the survey. All questions in this survey will focus on one particular feature, but we will ask about your opinion on how different people should be able to use it.

The following use the example “Your Spouse”, a “Smart Hub”, and a hub-related capability.

The questions on this page only focus on the following person: **Your spouse**: Imagine you have a spouse. You live with them everyday and share all smart appliances in your home. You make decisions together in most cases, especially important ones.

Imagine you are the owner of a **Smart Hub**.

Should **your spouse** be able to use the following feature? [**capability**]
 Always (24/7/365) Never Sometimes, depending on specific factors

Show questions if “Always” chosen
Why?

Imagine that the device incorrectly denies your spouse the ability to use this feature. How much of an inconvenience, if any, would this be? Not an inconvenience Minor inconvenience Major inconvenience

Why? Please be specific.

Show questions if “Never” chosen
Why?

Imagine that the device incorrectly allows your spouse the ability to use this feature. How much of an inconvenience, if any, would this be?

Not an inconvenience Minor inconvenience Major inconvenience

Why? Please be specific.

Show questions if “Sometimes” chosen

When should they be allowed to use this feature? Please be specific.

How important is it that they be allowed to use the feature in the cases you specified above? Not important Slightly important Moderately important Very important Extremely important

In contrast, when should they not be allowed to use this feature? Please be specific.

How important is it that they not be allowed to use the feature in the cases you specified above?

Not important Slightly important Moderately important Very important Extremely important

Thanks! We will now be asking you an additional set of questions. Imagine that you have already chosen settings specifying who can and cannot access a certain feature in your home. Think broadly about all types of people you might want to allow to control these devices; do not restrict yourself just to the relationships we have previously asked about.

Scenario: Imagine you are still the owner of a **Smart Hub**. You specify that certain people can access the following feature only sometimes: [**capability**]

Might the **location of the person relative to the device** (e. g., in the same room, not in the house, etc.) affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might the **location of the device in the house** (e. g., which room) affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might the **current state of the device** (e. g., whether it is on or off) affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might the **cost of performing that action** (e. g., cost of electricity or other monetary costs of carrying out that action) affect your decision on whether certain people can or cannot use this particular feature?

Yes No Not applicable

Briefly explain your response.

Might the **person's recent usage of the device** affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might the **age of the person** affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might **who else, if anyone, is currently at home** affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Might the **time of day** affect your decision on whether certain people can or cannot use this particular feature? Yes No Not applicable

Briefly explain your response.

Please list any other factors that might affect your decision on whether certain people can or cannot use the following feature: **[capability]**

Do you or anyone in your household own the following devices?

- Internet-connected lights? Yes No
- Internet-connected thermostat? Yes No
- Internet-connected voice assistant? Yes No
- Internet-connected lawn mower? Yes No
- Internet-connected security camera? Yes No
- Internet-connected door lock? Yes No

If answered yes to any of the above: Which specific devices (brand, model, etc.) do you own?

Please choose the answer that best applies:

Spouse: I'm currently living with such a person I'm not currently living with such a person, but I have previously I have never lived with such a person I prefer not to answer

Child in elementary school: I'm currently living with such a person I'm not currently living with such a person, but I have previously I have never lived with such a person I prefer not to answer

Teenage child: I'm currently living with such a person I'm not currently living with such a person, but I have previously I have never lived with such a person I prefer not to answer

Which of the following best describes your experience with hiring a babysitter (someone unrelated to you whom you pay to watch your children)? I have hired a babysitter within the last year I have hired a babysitter but not within the last year I have never hired a babysitter I prefer not to answer

Which of the following best describes your neighbors? I have neighbors and I know most of them I have neighbors and I know some of them I have neighbors and I know few or none of them I do not have neighbors I prefer not to answer

In a typical year, how many nights total do relatives (who do not live with you) stay at your home? 0 1-10 10-20 20-30 30+ I prefer not to answer

Do you live in a: Single family home Townhouse Apartment/condo Other (please specify) I prefer not to answer

Do you rent or own the place where you live? Rent Own I prefer not to answer

How many people (including you) are there in your household? 1 2 3 4 5 More than 5 I prefer not to answer

What is your age range? 18-24 25-34 35-44 45-54 55-64 65-74 75+ Prefer not to say

With what gender do you identify? Male Female Non-binary Other _____ Prefer not to say

Are you majoring in, hold a degree in, or have held a job in any of the following fields: computer science; computer engineering; information technology; or a related field? Yes No Prefer not to answer

If you have any further feedback, questions, comments, concerns, or anything else you want to tell us, please leave a comment below!