

How Risky Are Real Users’ IFTTT Applets?

Camille Cobb Milijana Surbatovich Anna Kawakami Mahmood Sharif
Carnegie Mellon University *Carnegie Mellon University* *Wellesley College* *NortonLifeLock*

Lujo Bauer Anupam Das Limin Jia
Carnegie Mellon University *North Carolina State University* *Carnegie Mellon University*

Abstract

Smart-home devices are becoming increasingly ubiquitous and interconnected with other devices and services, such as phones, fitness trackers, cars, and social media accounts. Built-in connections between these services are still emerging, but end-user-programming tools such as *If-This-Then-That* (IFTTT) have existed for almost a decade, allowing users to create rules (called *applets* in IFTTT) that dictate interactions between devices and services. Previous work found potential secrecy or integrity violations in many applets, but did so without examining how individual users interact with the service. In this work, we study the risks of real-world use of IFTTT by collecting and analyzing 732 applets installed by 28 participants and participants’ responses to several survey questions. We found that significantly fewer applets than previously thought pose realistic secrecy or integrity risks *to the users who install them*. Consistent with this finding, participants were generally not concerned about potential harms, even when these were explained to them. However, examining participants’ applets led us to identify several new types of privacy risks, which challenge some assumptions inherent in previous analyses that focus on secrecy and integrity risks. For example, we found that many applets involve monitoring *incidental users*: family, friends, and neighbors who may interact with someone else’s smart-home devices, possibly without realizing it. We discuss what our findings imply for automatically identifying potentially harmful applets.

1 Introduction

Smart home technology has made its way into public consciousness and widespread use [3]. On their own, smart-home devices typically allow users to control them via dedicated apps, possibly creating schedules, routines, or triggering notifications from the apps on users’ phones. Additionally, many smart-home devices enhance their capacity for home automation by interfacing with end-user programming tools such as If-This-Then-That (IFTTT), Stringify, and WebHooks. Such tools allow users to create trigger-action “rules” that react to and/or control their IoT devices and services like social media, cloud storage, or news. This enables users to accomplish home automation tasks that would not be possible otherwise. For example, a user could create a rule to automatically turn on all their smart lights when they arrive home, even if those lights were made by a variety of manufacturers. While these tools can enable creative, beneficial uses of smart-home technologies, they may also introduce security and privacy risks.

Prior work found that as many as 50% of applets shared on the IFTTT webpage could lead to secrecy or integrity violations (i.e., leak private information or allow unauthorized access to a user’s devices and services) [35]. That study, and others (e.g., [8, 10, 11, 28, 38]), sought to understand and measure the prevalence and magnitude of security and privacy risks of end-user programming with trigger-action rules, and they have proposed automated ways of identifying risky rules—rules that have the potential to cause harm—with an end-goal of mitigating risks. However, these studies have relied on publicly available data (e.g., applets shared on the IFTTT webpage) and have not evaluated risks in the context of individual users’ sets of rules, the contexts in which those rules are applied, or the individuals’ privacy preferences.

In this paper, we seek to better contextualize our understanding of the ways that users employ end-user programming in order to answer open questions about the secrecy, integrity, and other security and privacy risks their rules may create. To do so, we focus specifically on IFTTT, which is the most popular end-user-programming tool [25]. We recruited 28 IFTTT

users via popular home-automation message boards. Participants allowed us to collect data about their IFTTT applets and responded to a short survey. Survey questions addressed the context in which the applets are used (e.g., who cloud storage documents are shared with), participants’ understanding and perception of secrecy and integrity risks (e.g., if they had considered certain risks when setting up rules, if they had experienced any harms, and if they believed certain risks were possible for a particular rule), and how they would react to specific violations identified in prior work.

Using automated information-flow-based analysis, we found that about 57% of participants’ IFTTT rules had potential secrecy or integrity violations (see Section 4.3), which is consistent with the findings of prior work analyzing applets shared on the IFTTT website. In Section 4.4, we examine participants’ rules in more detail, considering context such as their titles. This more detailed analysis revealed that although many applets might technically have secrecy or integrity violations, they are rarely harmful because of these violations. Only about 10% of the secrecy-violating rules (just over 3% of all rules) could lead to secrecy harms, and just under 20% of integrity-violating rules (8.6% of all rules) present serious integrity-related risks. Consistent with our manual evaluation, participants did not believe that their rules were likely to lead to secrecy- or integrity-related harms, though they did care about the security and privacy of their rules.

Our contextualized analysis of trigger-action rules and their security and privacy risks is a key contribution of this work and also led to unexpected findings. Although secrecy and integrity violations rarely pose risks to IFTTT users, IFTTT rules pose other types of security and privacy risks that have not been identified through automated analysis. For example, IFTTT rules can create surveillance risks to *incidental users*—people besides the IFTTT user who created the rule. In Section 5, we discuss these other types of risks, as well as other limitations of the information-flow analysis. From our findings we draw guidelines for how automated analysis tools could better distinguish between practically *risky* and merely theoretically *violating* trigger-action rules. We also propose future research to better understand incidental users’ preferences regarding their interactions with smart-home devices. Identifying contextual factors needed for more accurate automated analyses and previously unexplored categories of risks are also key contributions of this study.

2 Background

2.1 Security of Smart-Home Technology

In recent years, researchers have investigated the security and privacy risks imposed by home IoT ecosystems. Most of these efforts consider the IoT ecosystem either at the application level or at the network level. At the application level, researchers have found that many applications built on emerging

programming platforms such as Samsung’s SmartThings [4] are over-privileged due to design flaws in their permission models [15, 17]. User-centric and context-aware permission systems have been developed for appified IoT platforms to address their coarse-grained permission flaws [16, 23, 37]. Systems utilizing static analysis [10, 28], model-checking [11], and data provenance graphs [38] have been proposed to help identify incorrect or inconsistent application behavior. Many research groups have proposed network-traffic-analysis-based security mechanisms [9, 12, 13, 29, 33, 34, 40]; many of these were introduced in light of the infamous Mirai attack, which took advantage of insecure IoT devices to launch a distributed denial of service (DDoS) attack [20, 30].

Differently from these studies, our work focuses on risks introduced by end-user programming. That is, we find that potential harms persist even under the assumption that technical vulnerabilities do not exist or are sufficiently unlikely.

2.2 Privacy Concerns in Smart Homes

In spite of their widespread adoption, users continue to surface privacy concerns about smart-home devices. To understand what concerns users have about smart-home technology, several interview- and survey-based studies investigated users’ experiences and preferences [6, 7, 14, 36]. When IoT devices are installed in multi-person households, new security, privacy, and usability challenges emerge. Recent research has sought to identify user requirements in these multi-user settings and proposed potential solutions [19, 39, 41] such as making it easier for everyone in a household to control the devices and how they are configured [41]. Others have studied desirable access controls for smart-home devices [21, 32]. Our study also attempts to understand privacy concerns in a smart-home setting (including multi-user setting), but more so in the context of using automation services like IFTTT which can inadvertently cause harms.

2.3 End-User Programming for IoT Devices

Several end-user programming tools—including IFTTT (“If This, Then That”) [1], Microsoft Flow [2] and Zapier [5]—enable users to connect multiple services by constructing simple trigger-action programs [24]; IFTTT is by far the most popular of these [25].

2.3.1 IFTTT

An IFTTT rule or “applet” (previously called “recipe”) consists of a “trigger” and an “action.” The trigger is the “this” and the action is the “that” in “if this then that.” Shortly before our study’s data collection, IFTTT added a feature to allow a single applet to have more than one action. Each trigger and action belongs to a “channel,” which specifies the service provider who created the trigger or action (e.g., IoT device

manufacturer, social media company). As of November 2019, IFTTT offered 1,228 channels [1]. Some actions and triggers have additional fields that must be specified by a user. For example, in the trigger “Amazon Alexa :: say a specific phrase,” a user-configured field specifies the specific phrase. When users set up rules, they can edit a plaintext description of the applet, which we refer to as the applet’s “title.”

2.3.2 Information-Flow Analysis of IFTTT Applets

Although our main focus is to understand the potential harms of real users’ IFTTT rules, we sought to ground our assessment in terms of previous estimates. Doing so enables us to assess the efficacy of previous methods of assessing IFTTT rules and to contribute insights that may improve automated analysis methods. In particular, we build on a prior study by Surbatovich et al. that applied information-flow analysis to IFTTT applets to automatically determine which rules contain potential secrecy or integrity violations [35].

Secrecy violations occur when a rule allows information to flow from a more private source (the trigger) to a less private sink (the action), possibly leaking private information to a wider audience than intended. For example, a rule that posts to Facebook each time motion is detected at the user’s front door could unintentionally broadcast when the user arrives home at a suspicious time (e.g., late at night or when they should be at work).

Integrity violations occur when a rule allows a more trusted action to be *controlled* by a less trusted trigger, thus possibly allowing unintended people to perform actions they would not otherwise be able to (e.g., allowing an adversary to control a user’s smart-home devices or post to their social media pages). For example, a rule that unlocks the user’s home when an email is received that contains a pre-specified keyword could allow an adversary who guesses the keyword to compromise the user’s home security.

Information-flow analysis as used in previous work [35] consists of three steps: (1) creating a set of secrecy and a set of integrity labels and arranging each set into a lattice that describes whether information flows are safe or constitute secrecy or integrity violations; (2) manually assigning *secrecy* and *integrity* labels to each trigger and action, conveying who can observe (secrecy) or control (integrity) the trigger or action; and (3) given the labels of trigger-action pairs and the secrecy and integrity lattices, determining whether the trigger-action pair constitutes a security or integrity violation.

Surbatovich et al. used four secrecy labels to describe who may have observed that a trigger or action event has occurred: *private* (e.g., only the IFTTT user), *restricted physical* (e.g., people in a user’s home), *restricted online* (e.g., the user’s Facebook friends), and *public* (e.g., anyone in the world). Similarly, they used six integrity labels to describe who can cause a trigger or action to occur: *trusted* (e.g., only the IFTTT user), *trusted other* (e.g., trusted news sources or weather re-

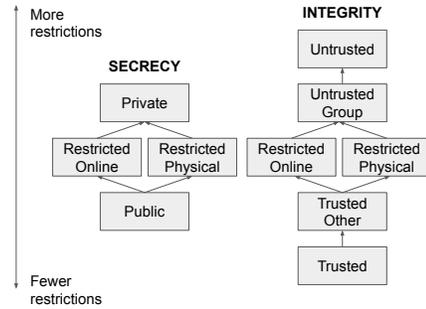


Figure 1: Secrecy and integrity lattices from prior work [35].

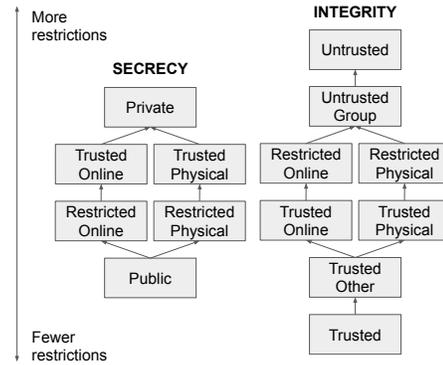


Figure 2: Modified secrecy and integrity lattices used in our analysis.

ports), *untrusted* (e.g., anyone in the world), *untrusted group* (e.g., the unrestricted group people could collectively cause a trigger event such as “reddit :: New hot post in subreddit”), and *restricted physical* and *restricted online* (same as for secrecy labels). Detailed explanations of the labels are included in Appendix B. We replicate this analysis with minor adjustments as described in Section 3.1.

When a precise secrecy or integrity label cannot be determined (e.g., because context, such with whom a Google Sheet is shared, would affect the label), multiple labels may be assigned to denote that the label could be *any* of these. If some combination of labels for an applet constitutes a secrecy or integrity violation, we refer to it as “potentially violating.”

2.3.3 Other Automated Analyses of Applets

Several other research efforts have sought to identify and characterize security and privacy risks for IFTTT users. Bastys et al. found that applets shared on IFTTT for others to use are a potential attack vector [8]. In user-shared rules, filters and parameters set by the rule creator are not visible to users who adopt the rule, but can be manipulated to execute URL attacks. Wang et al. and Hsu et al. considered how “chains” of trigger-action rules that cause other rules to execute can lead to significant complexity in evaluating whether a system of connected

smart-home devices is working as intended [22,38]. They proposed solutions that provide data provenance [38] or that use dynamic model-checking techniques to prevent hidden attack chains from executing [22]. IFTTT’s use of long-term OAuth tokens to connect various services creates another potential source of risk. To prevent the misuse of these tokens, Fernandes et al. proposed a Decentralized Trigger-Action Platform (DTAP) with fine-grained rule-specific tokens [18]. Our study differs from these in that we focus on analysis at the rule level; that is, we consider only risks that could arise while the applets themselves execute correctly as defined.

3 Method

Our study consisted of an online survey and automated collection of IFTTT data for 28 participants. Both were conducted with participant consent and approved by our institution’s IRB. We recruited participants via posts on Home Automation Reddit and SmartThings and Stringify forums, which advertised it as a study about “the habits and behavior of IFTTT users.” Participants received a \$5 Amazon gift card, and one out of ten participants chosen at random received an additional \$50 Amazon gift card. Participants were required to have an IFTTT account with at least five applets, in order to ensure that participants were active IFTTT users and because many survey questions referenced participants’ specific applets.

Of the 67 people who started our study, 41 completed it (i.e., 61% completion rate). The main contributor to dropout was people opting out of downloading a Chrome extension to collect their IFTTT data. We also excluded 13 responses from people who did not have the requisite five applets to qualify. Thus, our final data set included IFTTT data and survey responses from 28 participants. The survey was open for eight weeks in April and May of 2018. The median time to complete the study was 18 minutes (IQR 16-24 minutes). At the end of the survey, we revealed its purpose as a security-based survey to “see users’ perceptions and awareness of applets that are possibly insecure or safe,” and offered participants to opt out. No participants opted out at this stage.

3.1 IFTTT Data Collection and Analysis

We created a Chrome extension to collect the applets participants had installed or created, specifically, the trigger and trigger channel, the action(s) and action channel(s), and the applet title. Participants downloaded our extension and signed into their IFTTT account so that we could collect their applets; the extension was subsequently automatically removed.

Post-hoc data analysis had two phases: (1) automated information-flow analysis to identify applets with potential secrecy and privacy violations, and (2) a qualitative analysis of each trigger and action, to group similar triggers and actions.

3.1.1 Information-Flow Analysis

Surbatovich et al. shared with us their data and codebase [35], which allowed us to reproduce the assessment of secrecy or integrity violations as they would have been evaluated in that study, as described in Section 2.3.2. This allowed us to directly compare the secrecy and integrity violations in the recipes we collected with what was previously measured. Because many new triggers and actions were added to IFTTT after the earlier study, only about 20% of participants’ rules could be assessed with the pre-existing secrecy and integrity labels.

To compensate for this and for changes to the functionality of devices and services may have altered the appropriate label for existing triggers and actions (e.g., smart-home devices that added new online access capabilities for multiple household members), we relabelled *all* of the triggers and actions, rather than only adding labels for new triggers and actions. To relabel, four researchers met in three labeling sessions, totaling over six hours of meeting time. We looked up and discussed device functionality, which always allowed us to reach consensus. For example, the secrecy and integrity labels for a device can depend on whether the device’s default app supports just one or multiple users.

During relabelling, we found that the original information-flow analysis was limited in its ability to differentiate between those *restricted online* groups that could be used to grant family access to smart-home devices and the much larger *restricted online* groups such as all Facebook friends. The previous analysis also did not consider situations when two *different* restricted groups had simultaneous access (e.g., smart lights that can be controlled by people in a *restricted physical space and* by household members via an app).

To address these limitations, we made two adjustments to the security lattices used in prior work. First, we added two labels, both of which apply to both secrecy and integrity—*trusted online* and *trusted physical*—to represent very small, trusted groups of individuals. Second, we adapted the security lattice to include *unions* of two labels (as is standard in information-flow analyses [26,27]), so that we could indicate when, for example, a *restricted physical* and *trusted online* group could *both* control a device, as is the case in many smart-home devices with apps. The modified secrecy and integrity lattices are shown in Figure 2. The information-flow labels for all triggers and actions are included in Appendix C.

3.1.2 Semantic Labelling of Triggers and Actions

To evaluate whether applets that were potentially violating are likely to lead to harm, we identified *semantic labels* to group together similar triggers and actions. For example, 20 different actions control smart lights, and the distinction between a voice command given to Alexa versus Google Assistant is likely unimportant for evaluating risks. In terms of evaluating the riskiness of rules, these semantic labels allow us to examine what an attacker could learn in the event of a secrecy

violation; and what an attacker could control in the event of an integrity violation. A single researcher identified semantic labels via open coding after repeatedly and iteratively discussing trigger and action categories collectively with the other researchers.

Semantic labels for triggers are: *Weather or time*, *News-ish*, *Sensing IoT device state*, *Environment sensing*, *Intentional trigger*, *Voice command*, *Incoming communication*, *Sensing online account state*, *Actions with personal devices*, and *Other automations*. Semantic labels for actions are: *Change IoT device state* (with optional sub-labels of *Home security* and *Lights*), *Log or notify*, *Change personal device state*, *Outgoing communication*, and *Other automations*. We intentionally created distinct categories for particularly popular types of triggers and actions (e.g., giving *Voice commands* their own label despite their similarity to *Intentional triggers*, because they are so prevalent). The labels’ titles generally sufficiently describe their meaning, but more detailed explanations are included in Appendix B. Semantic labels for all triggers and actions are included in Appendix C.

Two researchers independently applied these labels to to all 160 triggers and 112 actions, with a high degree of agreement (Cohen’s $\kappa = 0.93$, calculated separately for triggers and actions). We subsequently discussed and came to a consensus about disagreements. Details about rules with each semantic label are shown in Tables 3 and Table 4. In Section 4.4, we use these semantic labels as the basis for evaluating the riskiness of participants’ rules, frequently referring to representative examples of the triggers and actions included in each set.

3.2 Participant Survey

Participants answered survey questions addressing three broad topics: (1) how they choose applets, (2) their beliefs and preferences about security and privacy properties of their applets, (3) harms they had experienced from using IFTTT. The full survey instrument is available in Appendix D.

After they downloaded the browser extension, we asked participants general questions about their use of IFTTT, such as how often and whether they prefer to create their own applets and if they ever turn on applets based on a friend or colleague’s recommendation. We also asked participants, on a 5-point Likert scale, whether they agreed or disagreed that they would be comfortable with friends, colleagues, or “anyone” knowing what applets they use (see Table 6).

Next, we asked each participant a series of questions about up to five randomly chosen applets that were violating according to the analysis from prior work [35]. For each applet, participants were prompted to consider four situations (or five situations for applets that involve a physical device) in which the applet might contribute to harm and rate whether this would make them very upset, slightly upset, or not upset. The situations were chosen to reflect potential harms identified in previous work [35] and are listed in Table 5. Near the end of

Age	25-34 (17), 35-44 (6), 45-54 (5)
Gender	Male (27), Female (1)
Education	High school graduate (2), bachelor’s degree (16), mboxassociate’s degree (2), professional degree (Master’s/PhD) (8)
Household size	1 (2), 2 (7), 3 (10), 4 (6), 5 (2), 7 (1)
Other household members	Family (24), housemates (1), other (1), live alone (2)
IUIPC Score (Avg.)	Overall 6.01±0.66 Control 6.25±0.66 Awareness 6.58±0.41 Collection 5.23±1.44

Table 1: Summary of participant demographics.

the survey, we explained the concept of secrecy and integrity violations (using lay language) and asked if considering this changed participants’ desire to keep using any of their applets.

Several additional questions, spread throughout the survey, asked participants whether they had experienced concerns or incidents related to their applets’ security and privacy. For example, we asked participants whether they ever experienced an incident in which an applet made them feel unsafe or that their privacy was violated and whether they had ever manually deleted anything that was posted automatically by an applet. Participants who answered affirmatively were asked to elaborate. We concluded the survey by asking participants to provide demographic information. The survey included additional questions that we do not discuss because they were not directly applicable to the specific focus of this work.

4 Results

4.1 Participant Characteristics

Participants were predominantly men and highly educated. Table 1 shows participant demographics. Most participants (24) lived with at least one other person, typically a family member. Details about participants’ living situations can help inform our risk analysis of their applets. For example, since only two participants live alone, we can assume that physical devices connected with their rules are likely to be seen and interacted with by other members of the household (e.g., ambiguity between the labels of *trusted physical* or *trusted in* the information-flow analysis could be resolved).

4.2 Characteristics of Participants’ Applets

Participants had a total of 732 applets. Each participant had between 5 and 66 applets (average=26, sd=20). Most applets had a single trigger with a single action; however, seven applets had a single trigger with multiple actions. Multi-action applets were a newly-added IFTTT feature at the time of the survey. For example, P24’s applet titled “*Google Home Find My Phone*” used a voice command to (1) set the Android device’s ringtone volume (presumably increasing phone volume so it was easier to find) and (2) receive a phone call. For

most analyses, we use a demultiplexed version of the data, where multi-action applets are treated as multiple rules with the same trigger and different actions. To differentiate, we refer to the demultiplexed version of data as “rules” rather than “applets.” There were 743 rules in our data set.

Frequency of Triggers and Actions. Participants’ rules included 68 unique trigger channels with 160 unique triggers and 64 unique action channels with 112 unique actions—only a small fraction of the total triggers and actions available from IFTTT. 396 unique combinations of triggers and actions were represented in participants’ rules. For reference, we include a list of the top ten channels, triggers, actions, and trigger-action pairs (i.e., those used by the most participants) in Appendix A.

Frequency of Trigger-Action Pairs. Participants used diverse trigger-action pairs, and every participant had at least one trigger-action pair that was unique to them. 63% of unique trigger-action pairs occurred only once (251 unique pairs, 34% of rules). When trigger-action combinations appeared multiple times, it was often because the same person used the same trigger-action pair repeatedly. For example, P13 had three rules with the same trigger (*Alexa :: Say a specific phrase*) and action (*Philips Hue :: Set a scene in a room*). User-specified titles differentiated these rules, in this case conveying what the specific phrase was: “trigger bright mode,” “trigger sleep mode,” or “trigger read mode.”

89 unique trigger-action pairs were used multiple times by only a single participant, making up 251 rules (22% of unique trigger-action pairs, 34% of rules). Only 56 trigger-action combinations (14% of unique trigger-action pairs, 32% of rules) were used by more than one person, and most of these were used more than once by some participants.

Cloud Storage Sharing Settings. Participants frequently had applets with triggers or actions that access or modify cloud storage files. For example, 51 rules updated Google Sheets or Google Drive and 17 updated Dropbox. Other cloud storage tools included Evernote and Day One. In our survey, we automatically identified 41 rules involving Google Sheets, Google Drive, or Dropbox, and asked the participants about the document’s or cloud storage space’s sharing settings. Participants unanimously stated that the cloud storage space or document was not shared with *anyone* else. In Section 4.4.1, we discuss the implication of these cloud storage files being private on our assessment of the potential risks of these rules.

4.3 Information-Flow-Based Analysis

Table 2 shows the breakdown of rules and unique trigger-action pairs that were found to be violating using the secrecy and integrity labels from prior work [35], and the new labels and lattice we generated for this study. Exactly replicating the analysis from previous work to evaluate the 158 fully-labelled rules, we found that about 35% of rules were potentially violating, compared to about 50% in previous work [35]. Using the updated lattice and labels (see Section 3.1), applied to all

	Total rules	Potentially violating	Secrecy-violating	Integrity-violating
2017 lattice	158	55 (34.8%)	41 (25.9%)	48 (30.4%)
Updated lattice	743	423 (56.9%)	269 (36.2%)	319 (42.9%)

Table 2: Secrecy and integrity violations reported by the information-flow-based analysis.

743 rules, we found that 57% of rules were potentially violating. In total, 319 rules have potential integrity violations and 269 rules have potential secrecy violations.

However, a key research question remains: are the rules labelled as violating actually something to worry about and are the remaining rules actually innocuous?

4.4 Secrecy and Integrity Risks in Context

In this section, we assess the riskiness in practice of the rules deemed potentially violating by information-flow-based analysis. We use semantic labels (described in Section 3.1.2) to structure this analysis and leverage contextual information such as applets’ titles to better understand specific rules.

4.4.1 Evaluating Secrecy-Violating Rules

Out of the 269 rules with potential secrecy violations, almost one third are unlikely to actually carry sensitive information, and 42% (some of which overlap with that third) have actions that are probably not observable by unintended people. We judge these rules to be unlikely to be harmful based on the semantic labels assigned to either their triggers or their actions. We next describe why rules with specific semantic labels are unlikely to be harmful. In total, only around 10% of secrecy violating rules (a subset of those with actions that have semantic labels of *Outgoing communication* or *Other automations*) are likely to lead to significant secrecy risks.

No Secret Information. 47 rules of the 164 rules with *Voice command* triggers and 24 of the rules with *Intentional triggers* (totaling 27% of rules with potential secrecy violations) have potential secrecy violations according to information-flow-based analysis. This is because the *Voice command* or *Intentional trigger* could be done privately, and they result in an action with wider observability. For example, 10 of these rules control *Lights* (as their action), which might be observed by neighbors outside of the user’s home, potentially revealing when a person is home or which rooms they are using. These rules act as an alternative light switch and do not introduce additional secrecy risks beyond those of a normal light switch. Also, because the user actively decides to cause this action each time the rule executes, they can evaluate at time-of-use whether the potential secrecy leak is a problem.

Twenty one *News-ish* rules were found to have potential secrecy violations (8% of secrecy-violating rules). Although many *News-ish* triggers have a secrecy label of *public* and

can, therefore, not lead to secrecy violations, these 21 rules have triggers that require additional context to determine whether they utilize public or restricted information. For example, the trigger “*Twitter :: New tweet by a specific user*” should have a secrecy label of *restricted online* if the specific user’s Twitter account is set to private and *public* if not. P20’s rule “*Save every tweet from the US President*” triggers based on public tweets and is therefore not secrecy-violating. We manually evaluated the titles for these 21 secrecy-violating *News-ish* rules, which were not considered in the analysis from previous work or the automated information-flow-based analysis; based on their titles, we determined that 13 are not actually violating (4.9% of all secrecy-violating rules).

Action Remains Private. 113 rules (42% of secrecy-violating rules) have an action with the semantic label *Log or notify*. This includes rules that update Google spreadsheets (45 rules) or Dropbox files or folders (17 rules), add calendar events (15 rules), or send the user a notification (24 rules). It is possible that these actions could leak information to untrusted parties who have access to the documents, folders, or calendars, or can observe when the user receives a notification. Based on participants’ survey responses, we know that most of participants’ applet-connected Google Drive and Dropbox files were not shared with other people. Hence, rules connected to private cloud storage spaces are not actually secrecy-violating (based on how cloud storage content is shared) although they could be (if the cloud storage content was widely shared), and hence were detected as potentially violating secrecy by the information-flow-based analysis.

Notifications have a *restricted physical* secrecy label, because there are many situations in which someone else might see a user’s phone screen (e.g., if their phone is being used for navigation in a car or to play music at a party). Rules that have actions that send a notification could potentially leak information if their action is private. Although this is a real risk, most users are routinely exposed to this risk, even without using IFTTT—many smart phone apps such as email and SMS have (by default) the side-effect of showing a notification, often with a message preview. Thus, rules that have an action that sends a notification directly are no more risky than those that send the user an email or an SMS.

Secrecy Risk Is Limited by the Expressivity of the Action. For 76 rules (28% of secrecy-violating rules), if they leaked sensitive information, they would do so by *Changing the State of a Personal Device* or *Changing the state of an IoT device* (e.g., 21 secrecy-violating rules control *Lights* and 4 change the user’s phone volume). The extent to which these rules could leak private information is limited by the expressivity of these devices—many have only an “on” or “off” option. Additionally, there are a plethora of *other* ways the action could occur (e.g., triggered by another rule, controlled directly from the device’s dedicated app, or through physical interaction with the device). Private information *could* be leaked via these rules, but the risk is typically low.

Sending “Outgoing” Communication. 37 potentially secrecy-violating rules (14%) share information via *Outgoing communication* (i.e., social media, SMS or email), which could leak sensitive or private information to other people. As we discuss again in terms of its implications for rules with potential integrity violations, participants regularly use *Outgoing communication* actions to send information to themselves. For example, P20 has a rule called “*Receive an email diagnosis from Dash if your car experiences an issue*” that uses the action “*Gmail :: Send an email*” (rather than the action “*Email :: Send me an email*”). In several cases, even when the rule’s title does not specifically state that the outgoing content is sent to the user themselves, we can infer that this is the case. For example, 14 of these 37 rules belong to P28 and post to Slack based on *Sensing IoT device states* (e.g., “*If basement Sliding Door closed then post a message to a Slack service*”); these rules probably post to a private Slack channel.

There are only 14 secrecy-violating rules (5%) which are likely to *actually* send *Outgoing communication*, which could be risky. For example, because of the rule “*Tumblr Likes to Pinterest*,” P9 could accidentally “like” a post on Tumblr that would be embarrassing if it was sent to his Pinterest followers. These secrecy risks exist even for rules with triggers and actions that *both* use the *same* channel, or service. For example, P18’s rule “*Save Facebook photos you’re tagged in to your own album*” could result in unflattering photos of him being added to his album, possibly with a broader audience on Facebook than the original post.

Other Secrecy-Violating Rules. Of all 269 rules found to potentially violate secrecy, the previous discussion has addressed all but 11 rules (4%). These 11 rules all have actions that are *Other [non-IFTTT] automations*. In particular, one rule has both a trigger and an action that are *Other automations* (“*If maker Event ‘mancave sleep’ then run a Stringify Flow*” by P13). Without additional details about the automations, we cannot evaluate the potential harms of these rules. It is probably pertinent to warn users who install such applets that no automated analysis could evaluate secrecy properties, which could make these rules especially risky.

4.4.2 Evaluating Integrity-Violating Rules

Although 319 rules have potential integrity violations, according to information-flow analysis, very few of them are actually likely to lead to integrity-related harms. 64% of the integrity-violating rules do nothing more than update a digital log or notify the user, which is unlikely to be harmful even if it is caused by an adversary. The rules that are most likely to be potentially risky include 40 rules with *Other automation* actions and 21 rules that would potentially allow an adversary to control smart home devices other than lights—totaling just under 20% of the integrity-violating rules.

Trigger Is Sufficiently Trusted. Many rules have triggers with *trusted other* integrity labels and actions labeled *trusted*.

While such rules are flagged as integrity violations, we found that rules with *trusted other* triggers would not typically create integrity risks. For example, the trigger “*Best Buy :: product price changes*” is controlled by a company (Best Buy) that is unlikely to change the price of a product with the goal of adversarially triggering someone’s applet. Out of 319 rules that have potentially integrity violations, 71 rules (23%) have similar triggers to the one discussed above, with *trusted other* integrity. These predominantly have the semantic labels *Weather or time* (30 rules) or *News-ish* (29 rules). The triggers “*Nest Protect :: Battery is low*” and “*Fitbit :: Daily activity summary*,” which is sent at the same time each day regardless of the activity summary’s contents, account for the remaining 12 rules with other semantic labels.

An additional 31 integrity-violating rules (10%) have an action with the semantic label *News-ish*, but not *trusted other* integrity. Based on manual examination of the applet titles, we determined that for 26 of these 31 rules the trigger has *trusted other* integrity in practice (8% of integrity-violating rules). For example, a new item in an RSS feed *could* come from an *untrusted* source; however, in P20’s rule “*Text me if the CDC reports a zombie outbreak*,” if the update comes from the United States Centers for Disease Control (CDC), as suggested by the title, then *trusted other* would be a more appropriate integrity label for this trigger.

Creating a Log or Notifying the User. The majority of integrity-violating rules (64%; 204 rules) have an action with the semantic label *Log or notify*. Some of these violations could sometimes lead to practical harm. For example, as noted in prior work, an adversary could potentially fill up cloud storage space [35], or an ill-timed notification could disrupt an important meeting. However, their titles reveal that many of these applets *intentionally* trigger based on other people’s actions. For example, in P3’s rule “*If office Nest Protect battery is low, then send a notification*,” other people with physical access to the home *could* cause this rule to execute (e.g., by repeatedly touching the device to keep the screen on and drain the battery more quickly). The user would likely still want this warning so that they know to replace the battery.

Sending “Outgoing” Communication. 27 rules (8% of integrity-violating rules) have an action that sends *Outgoing communication*. Someone who controls the trigger could, for example, spam the user’s friends with emails or create unwanted social media posts, if that is what the applets were set up to do. In practice, however, as previously discussed regarding secrecy-violating rules, many of these do not actually send *outgoing* messages. The titles of these 27 rules reveal that almost all of them likely send information only to the IFTTT user. For example, P10’s rule “*Have Alexa email you your shopping list*” probably emails P10, despite using the action “*Gmail :: send an email*” (*Outgoing communication*) rather than “*Email :: send me an email*” (*Log or notify*).

Controlling Smart-Home Devices. 37 rules with potential integrity violations (12%) have an action that *Changes an*

IoT device’s state, including 16 that control *Lights* and 9 that control devices related to *Home security*. If these rules were triggered maliciously, the extent and type of harm, is predominantly determined by the capabilities of the device they control and other contextual factors. An additional 41 integrity-violating rules (13%) have an action that causes a non-IFTTT automation to execute. For 21 of these 41 rules, the titles suggest that they change the state of home IoT devices (e.g., P27’s rule “*If You say ‘Set Sonos to 10 percent then run a Stringify Flow’*”).

Rules that control lights, could be harmful if they are used at inopportune times—e.g., lights coming on at 2am causing someone to lose sleep, lights turning off while someone is walking down stairs, potentially causing an injury. A more likely risk is that lights are left on more than they normally would be, consuming electricity or causing the device to wear out more quickly. This creates a financial risk bounded by the cost of the device plus the cost of the light being on constantly. Other types of smart-home devices might have a greater potential to cause expensive and/or dangerous damage. For example, P7 and P12 have rules that mention turning on a waffle iron and “cat heaters,” respectively, which could potentially start a fire. P9’s rules that turn on an irrigation systems cause costly damage if used during freezing weather.

Users may be especially protective of rules that affect their home security (40 total rules, 9 of which have potential integrity violations). For example, P27’s rule “*If You say ‘Disarm Blink’ [to Google Assistant] then disarm Outside Blink system*” is a rule that might warrant a warning to the user; an unintended person could potentially speak loudly enough from outside the home to disarm the system. However, many users might still decide that this risk is acceptable or sufficiently unlikely given the placement of their smart assistant.

Controlling the User’s Personal Device. 10 integrity-violating rules (3%) *Change [the user’s] personal device state* (e.g., changing the volume level or launching an app like music or navigation). If properly timed, an adversary could cause this rule to execute during an important meeting, causing embarrassment or punishment. However, similar risks exist any time a user’s phone is on, unrelated to IFTTT (e.g., repeatedly calling during an important meeting). Alternately, an adversary could cause these rules to execute with the goal of draining the user’s phone battery more quickly than usual, which could be dangerous in some contexts (e.g., if the user is in an unfamiliar place and will need directions or a ride home). Launching certain apps could utilize cell data, which might be limited or expensive for the user.

Other Integrity-Violating Rules. 40 integrity-violating rules (13%) have not yet been addressed. All of these rules have actions that are *Other [non-IFTTT] automations*. As with secrecy-violating rules that have *Other automation* actions, we cannot evaluate the possible integrity harms that could be associated with these rules. Therefore, these rules should be treated with caution.

Semantic trigger label	# (%) of rules	# potentially violating rules (secrecy / integrity)
<i>Weather or time</i>	101 (13.6%)	30 (0 / 30)
<i>News-ish</i>	121 (16.3%)	60 (21 / 60)
<i>Sensing IoT device state</i>	111 (14.9%)	68 (45 / 57)
<i>Environment sensing</i>	46 (6.2%)	40 (7 / 37)
<i>Intentional trigger</i>	38 (5.1%)	24 (24 / 0)
<i>Voice command</i>	164 (22.1%)	72 (47 / 72)
<i>Incoming communication</i>	19 (2.6%)	15 (17 / 14)
<i>Sensing online account state</i>	57 (7.7%)	49 (46 / 20)
<i>Actions with personal devices</i>	58 (7.5%)	33 (33 / 0)
<i>Other automations</i>	30 (4.0%)	29 (29 / 29)

Table 3: The distribution of rules with each of five semantic trigger labels, and the breakdown of rules with each label with secrecy or integrity violations.

Semantic action label	# (%) of rules	# potentially violating rules (secrecy / integrity)
<i>Change IoT device state</i>	254 (34.2%)	79 (72 / 37)
<i>Home security*</i>	40 (5.4%)	17 (14 / 9)
<i>Lights*</i>	66 (8.9%)	25 (21 / 16)
<i>Log or notify</i>	368 (49.5%)	246 (113 / 204)
<i>Change personal device state</i>	28 (3.8%)	11 (4 / 10)
<i>Outgoing communication</i>	41 (5.5%)	39 (37 / 27)
<i>Other automations</i>	52 (7.0%)	49 (43 / 41)

Table 4: The distribution of rules with each of five semantic action labels, and the breakdown of rules with each label with secrecy or integrity violations. Labels denoted with an asterisk (*) are secondary labels for *Change IoT device state*.

4.5 Survey Responses

We now consider participants’ responses to survey questions, to help further contextualize our findings about their applets. Are participants’ assessments of and experiences with their own applets consistent with our finding that most applets are unlikely to lead to harm due to secrecy or integrity violations? What harms have participants encountered, including but not limited to those that arise because of secrecy or integrity violations?

4.5.1 Choosing Applets

Prior work hypothesized that applets shared publicly on IFTTT are a potential attack vector [8]. Most participants (16, 57%) reported a preference to create their own applets. 25 participants (89%) reported creating *some* applets themselves, and seven participants had created 20 or more applets. Thus, although preventing malicious applets from being available on the IFTTT webpage can mitigate some security and privacy risks, it is also important to be able to identify potential risks in applets created by users themselves.

In addition to using existing applets shared on IFTTT, 9 participants (32%) said that they sometimes or often turn on applets based on friends’ or colleagues’ recommendations. When taking suggestions from trusted sources, users might

Would you be upset if this applet contributed to the following situation occurring:	Very upset	Slightly upset	Not upset
TOTAL	105 (41.7%)	114 (45.2%)	33 (13.1%)
You no longer directly control what files are downloaded from email or social media, possibly spreading malware on your computer?	32 (69.6%)	9 (19.6%)	5 (10.9%)
Your electronic device is used in a way it wasn’t designed for (such as being toggled on/off very rapidly), possibly reducing its longevity or damaging it? *	21 (55.3%)	16 (42.1%)	1 (2.6%)
Private information gets posted online unintentionally, possibly embarrassing you?	28 (52.8%)	17 (32.1%)	8 (15.1%)
Data gets uploaded to your cloud storage more often than you thought, possibly causing you to run out of space?	14 (26.9%)	26 (50%)	12 (23.1%)
You consume more resources (e.g., electricity, phone data, cloud storage space), possibly increasing your bills or otherwise causing you to spend more money?	10 (15.9%)	46 (73%)	7 (11.1%)

Table 5: Participants typically stated that they would be very or slightly upset if an applet contributed to a harmful situation. The asterisk (*) denotes that this question applied only to rules that utilize a physical device in their trigger or action.

be less likely to consider the potential harms.

4.5.2 Participants Believe Their Applets Are Safe

In general, participants did not express concerns about security and privacy risks arising from *their own* use of IFTTT, though they seemed to be aware of the possibility that applets could lead to security and privacy risks. This is consistent with our assessment of their applets. Almost all participants (96%) believed that their applets work as expected and are safe to use. Only six participants (21%) changed their views on the riskiness of applets, even after we explained the definition of secrecy and integrity violations and how this might manifest in applets (at the end of the survey). All six reported increased caution about their applets’ secrecy; P16 also reported increased concern about integrity violations.

4.5.3 Harms Experienced from IFTTT Rules

Despite an overall sense that their applets are safe, four participants noted that they *had* experienced harms that they attributed to their applets or expressed that their applets sometimes did not work as expected. The harms they described were not the result of secrecy or integrity violations. Rather, they described instances in which the app or service malfunctioned or in which they had misconfigured a rule. P10’s door was unlocked when location updates functioned in an unexpected way: “... *I was in Disneyland and had just turned on my phone’s location settings, so when the IFTTT app received the first location broadcast, it was before GPS and cell location could be locked in and so it assumed I was still at home.*” P26 explained that she “*used to use an applet that posted*

my door lock to a google calendar. That got annoying.” P23 recalled an applet that accidentally flooded his Twitter with posts: “[it] was only supposed to trigger in some situations, I set it up wrong, realized it was posting too often.” Illustrating the complexity that can exist in a set of rules, P20 reported that one of his applets created an undesired Facebook post, though he was “unsure of which applet did it.” In Section 5, we discuss other types of harm, including but not limited to the ones these participants experienced.

Selection bias may have contributed to our finding that so few participants experienced harms due to their applets. Users who fear or experience harms due to their applets may be less likely to participate in online message boards about home automation, where we reached potential participants.

4.5.4 Participants Value Applets’ Security and Privacy

Although they mostly believed their applets were safe and did not change their level of caution based on explanations of potential secrecy and integrity violations, participants conveyed that the security and privacy of their applets *is* important to them. 21 participants (75%) said they would be upset if an applet triggered when they did not intend it to, which could happen through an integrity violation (or misconfiguration or incorrect behavior of the rule or connected services).

When asked about whether they would be upset if a specified, potentially-violating applet contributed to one of five undesirable outcomes, they only reported that they would not be upset for a total of 13% of applets (see Table 5). Comparing across the different types of harmful outcomes, they were less concerned about the possibility of using up cloud storage space and consuming extra resources than they were about an applet possibly posting private information that would embarrass them, spreading malware on their computer, or damaging their physical smart-home devices.

Many participants were uncomfortable with certain other people knowing which applets they have, especially strangers (i.e., 46% of participants disagreed that “[they are] comfortable with anyone knowing what applets [they] use,” as shown in Table 6). strangers having this information (see Table 6). In Section 5, we re-examine the implications of this finding.

5 Discussion

In contrast with the results of the automated analysis (Section 4.3) and the findings of previous work [35], our analysis of real users’ applets suggests that the majority of applets are unlikely to lead to significant risks due to secrecy or integrity violations. Nevertheless, our finding that participants *are* concerned about the security and privacy of their applets emphasizes the importance of identifying potentially harmful applets more accurately.

With this in mind, we next consider what *additional* harms applets could introduce that may not be encompassed by the

Survey Prompt	Strongly or somewhat agree	Neither agree nor disagree	Strongly or somewhat disagree
I am comfortable telling my friends what applets I use.	24 (85.7%)	3 (10.7%)	1 (3.6%)
I am comfortable telling my colleagues what applets I use.	23 (82.1%)	4 (14.3%)	1 (3.6%)
I am comfortable with anyone knowing what applets I use.	11 (39.3%)	4 (14.3%)	13 (46.5%)

Table 6: Participants were more likely to report that they would be comfortable with friends or colleagues knowing about their IFTTT rules than strangers. The answer choice with the most responses is shown in bold.

previous focus on secrecy or integrity risks, as they were defined in the automated analysis that we leveraged. We additionally re-surface key take-aways from Section 4, addressing limitations of the existing automated analysis that lead to non-risky applets being labeled as potentially violating. Finally, we conclude with guidelines for future tools that seeks to more accurately identify risky (or non-risky) applets.

5.1 Risks Beyond Secrecy and Integrity Violations Against the IFTTT User

Mismatch Between Reality and Expectations When Setting up an Applet. Several participants described scenarios in which their experiences using IFTTT applets did not match their expectations. This could occur for many reasons, which do not necessarily involve secrecy or integrity violations or adversarial actions: IFTTT services might not work as expected (e.g., not sensing location appropriately when devices connect and disconnect from the Internet), users might unintentionally misconfigure their applets (e.g., mis-spelling an applet parameter such as a search term), or users might not anticipate the impact of an applet even when it is configured as intended (e.g., P26 in retrospect decided that adding a calendar entry every time her door locked was “annoying”).

Surveillance Risks to Incidental Users. Many rules cause data to be collected about people *other* than the IFTTT user, possibly without their awareness. We refer to these other people as *incidental users*. For example, P11’s rule “If [name] presence detected, then create Journal entry” in effect monitors [name]’s location and schedule. Furthermore, all IFTTT rules create a record, accessible after the fact, whenever they execute. Hence, P9’s rule “If - Front Door locked then switch off Entryway Light” creates a record, through IFTTT, of each time the front door is locked. P9 could use this log to determine exactly when his family members arrive or leave.

Rules especially likely to trigger based on the actions of incidental users include those that *Sense IoT device states* (e.g., when a door is locked) or *Sense changes to the environment* (e.g., if the temperature increases when someone is home), as

well as those that trigger based on *Incoming communication*. Collectively, these make up 162 rules, or 22% of participants' rules, and include 53 rules that are non-violating according to information-flow analysis. Additionally, some rules with other semantic trigger labels might also present risks to incidental users. For example, several of P2's rules toggle a SmartThings Switch when three distinct people's phones connect to or disconnect from his WiFi (semantic trigger label *Actions with personal devices*). One of the phone owners is referred to as a "[baby/pet]sitter" in the applet title.

Incidental users can be family members, household visitors, employees, or neighbors. For example, P28's "Motion alert!" rule (which sends a notification when motion is detected) could inadvertently capture information about neighbors' daily schedules if it reacts to motion on a communal sidewalk. People other than household members might be especially unlikely to realize that they are being monitored via these rules. We have no reason to believe that any participants in our study used IFTTT in abusive ways to intentionally collect data about incidental users, but risks and harms, such as chilling effects [31], can exist even when surveillance is not intended adversarially. When such data is collected with malicious intent, as it might be in an abusive relationship, this type of surveillance could be especially harmful.

Our study was not designed to explore research questions related to incidental users and, thus, our conclusions on this topic are limited. Instead, our findings advocate for research that more thoroughly explores the privacy preferences and experiences of incidental users. Future work in this area would benefit from data collection of a larger scale and broader scope, including input from a variety of stakeholders other than people who use IFTTT or own smart home devices.

5.2 Limitations of Current Information-Flow-Based Automated Analysis

Trustworthiness of Information. Rules that convey trusted information could be used to trick users into trusting attacker-supplied information, particularly when this information is conveyed to the user over a less trusted channel. Several users, for example, used rules triggered by official CDC or weather information updates whose actions propagated this information to the user via email or another similar channel. In these cases, the source of the information is trusted, but the delivery channel used by the rule is not. An attacker could easily create an email message indistinguishable from one created by the rule, thereby tricking a user who set up the rule into believing that a trusted source supplied the information (and hence caused the email to be sent).

Such risks are not discovered by previous information-flow analyses: these would flag as a violation a flow from a lower integrity trigger to a higher integrity action; here, in contrast, the potential danger comes from the *user* treating an action as if it was as trustworthy as the trigger. Although this risk could

apply to any rule with a somewhat trusted trigger, *Log or notify* (193 rules) are explicitly designed to convey information and may be particularly susceptible.

Reconsidering What Is "Secret." The automated information-flow-based analysis we use assumes that IFTTT rules adopted by users are public (i.e., known to the attacker). Similarly, it assumes that secret information would flow from the trigger to the action. However, viewing real users' rules has led us to challenge these assumptions. What risks are substantially mitigated, or introduced, if we assume that adversaries do not know a rule and its configuration? Could adversaries infer sufficient details about rules to take advantage of otherwise unlikely risks?

Many potential secrecy and integrity violations in participants' rules could only be exploited if the adversary knew the rule. For example, P27 had the rule "*If You say 'Disarm Blink' then disarm Outside Blink [alarm] system.*" Without knowing the phrase to use to disarm the alarm system, an adversary may technically be able to disarm the alarm system but would have difficulty doing so. Similarly, without knowing about P11's rule "*If daily Steps goal not achieved by 10:15 pm, then send me an SMS,*" an adversary who sees (or hears) that P11 receives an SMS at 10:15pm would be unlikely to guess the *meaning* of the SMS (i.e., that P11 had been less active that day) without knowing that he was using this rule.

There are many potential ways that an adversary could infer details about the configuration of a rule (e.g., overhearing P27 use the smart assistant phrase to disarm their alarm system), and there are a variety of aspects of the rule that a user might care about keeping secret. In some cases, the secret or sensitive information may be entirely contained in the rule and only *implied* by the triggers or actions. For example, in the case of P11's rule "*If new [Craigslist] post from search Search URL then send me an SMS at [number],*" the user may wish to protect the secrecy of the specific search term they are following. The trigger is not secret, since Craigslist posts are publicly available; observations of the action could reveal this potentially sensitive detail about the rule.

By assuming that the IFTTT rule was public, the information-flow-based analysis on the one hand failed to identify rules that potentially leak sensitive information regarding the rule itself and on the other hand overestimated the probability of secrecy or integrity violations of other rules. Modeling the components of IFTTT rules on a more fine-grained level (e.g., specifying a secrecy and integrity label for the rule's parameters or the rule itself) could potentially address these limitations.

Rethinking the Granularity of Labels. In participants' applets, we found triggers and actions whose labels were both too fine grained and too coarse grained. For example, triggers based on official weather reports or the time of day were labelled as *trusted other*. This led to innocuous applets such as P17's "*Get the weather forecast every day at 7:00 AM*" being marked as potentially violating. In practice, distinguishing

between *trusted* and *trusted other* was unnecessary.

On the other hand, for some rules that have the same secrecy or integrity label for both their trigger and action—which causes information-flow analysis to judge them as safe—contextual details are needed to determine whether the people who can observe or control the trigger are in fact the same as those who can observe or control the action. For example, in P2’s rule “*If kitchen Lights switched off then turn off lights in Kitchen*” both the trigger and the action have *restricted physical* secrecy, so this rule is considered non-secrecy violating. Based on the title, we can infer that this interpretation is correct—the trigger and action occur in *same* physical space (the kitchen). A similar rule that triggers when bedroom lights are switched off would also be marked as non-secrecy violating; however, that rule might reveal to people in the kitchen when someone enters the bedroom, and hence does constitute a secrecy violation (albeit perhaps not a particularly harmful one). 224 (125) such non-violating rules have triggers and actions with the same secrecy (integrity) label (*restricted online*, *restricted physical*, *trusted online*, or *trusted physical*).

In order to more accurately identify when rules are actually violating, choices of secrecy and integrity labels should be better-informed by a deeper understanding of contextual factors (e.g., devices’ relative locations).

Challenges in Labelling External Services. Many services that access or update online content include attribution: they capture *who* added or edited content. The information-flow analysis specifies only that integrity labels describe “who could cause the event” [35]. We found that determining an integrity label for this type of action requires both a nuanced definition of “the event” and a deep understanding of (e.g., Google Calendar’s) functionality. For example, the integrity label for the action “*Google Calendar :: Quick add event*” should denote “who could add an event to this calendar” (*definition 1*). However, Google Calendar events include a “created by” field, specifying the user (account) that created the event. Since the IFTTT rule creates a calendar event that shows up as “created by [username]” (where [username] is the IFTTT user’s Google account name), an alternative integrity label would denote “who could add a calendar event to this calendar that appears to be created by [username]” (*definition 2*).

Although definition 1 is consistent with the action’s description in IFTTT, using it for information-flow analysis will not capture erroneous attribution of the calendar entry. If someone who can add events to a shared calendar (from their own user account) instead uses, e.g., Avery’s rule to create a calendar event, they could add events with embarrassing or offensive titles that other people would attribute to Avery.

5.3 Guidelines for Automatically Identifying Risky Applets

Based on our findings, we compiled guidelines to help more accurately identify potentially risky trigger-action rules:

- Be aware of gaps between users’ intent and installed rules (Section 5.1).
- Consider risks to incidental users, and consider a variety of potential adversaries such as abusive partners, which might include the IFTTT user themselves (Section 5.1).
- Analysis should not assume that other people (e.g., potential attackers) know the rule or its configuration (Section 5.2).
- Expect that appropriate secrecy and integrity labels are sensitive to contextual details that may be difficult to determine automatically, such as the settings of users’ external services or the location of their physical devices (Section 5.2).
- For giving intuitive warnings to users, semantic labels may be more useful than fine-grained analyses, because they correlate with risky rules (e.g., about *Home security*) and can explain how a rule could lead to harm. Simple heuristics could effectively complement more theoretically grounded analyses (Section 4.4).

Future work should seek to incorporate these guidelines into automated analyses and to provide a deeper understanding of users’ experiences. Our survey was designed to inquire about the types of harms posited in prior work; future work should seek to explore the other types of potential harm that we identified. For example, while we identified the possibility of IFTTT rules with surveillance risks to incidental users, we do not yet know if they recognize that these risks exist or what their security and privacy preferences regarding other peoples’ rules are. Our participant sample was demographically skewed; future studies could assess whether different groups of IFTTT users might have rules with different potential risks. Since incorporating contextual details may be necessary to determine appropriate secrecy and integrity labels, future work might consider how this information could be obtained, either by asking users directly or through automated (technical) means.

6 Conclusion

We evaluated the possible risks and harms associated with real users’ IFTTT applets. Applets were less risky than was previously shown through automated analysis that sought to identify secrecy and integrity violations; however, we discovered new types of potential harm not previously considered in that automated analysis. Additionally, we outline some of the ways that the automated analysis falls short even in its ability to accurately identify secrecy and integrity risks. Finally, we discuss guidelines for creating a better tool (future work) that would identify risky applets—both from the standpoint of more accurately identifying secrecy and integrity violations and in terms of identifying other types of risk.

Acknowledgments

This work was supported in part by gifts from Google and the CyLab Security and Privacy Institute at Carnegie Mellon University; by a CyLab Presidential Fellowship and a Symantec Research Lab fellowship; and by DARPA and the Air Force Research Laboratory under agreement number FA8750-15-2-0277.

References

- [1] If This Then That (IFTTT). <https://ifttt.com>.
- [2] Microsoft Flow. <https://flow.microsoft.com/en-us/>.
- [3] Research: 90 <https://www.iottechnews.com/news/2018/may/15/research-us-consumers-smart-home-device/>.
- [4] SmartThings. <https://www.smartthings.com/>.
- [5] Zapier. <https://zapier.com/>.
- [6] Noura Abdi, Kopo M. Ramokapane, and Jose M. Such. More than smart speakers: Security and privacy perceptions of smart home personal assistants. In *Proceedings of the 15th Symposium on Usable Privacy and Security (SOUPS)*, 2019.
- [7] Noah Aporthe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. Discovering smart home internet of things privacy norms using contextual integrity. *Proceedings of ACM Interaction Mobile Wearable Ubiquitous Technology*, 2(2):59:1–59:23, 2018.
- [8] Iulia Bastys, Musard Balliu, and Andrei Sabelfeld. If this then what? Controlling flows in IoT apps. In *Proceedings of the 25th ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1102–1119, 2018.
- [9] Suman S. Bhunia and Mohan Gurusamy. Dynamic attack detection and mitigation in IoT using sdn. In *Proceedings of the 27th IEEE International Telecommunication Networks and Applications Conference (ITNAC)*, pages 1–6, 2017.
- [10] Z. Berkay Celik, Patrick McDaniel, and Gang Tan. Soteria: Automated IoT safety and security analysis. In *Proceedings of the 2018 USENIX Annual Technical Conference (USENIX ATC)*, pages 147–158, 2018.
- [11] Z. Berkay Celik, Gang Tan, and Patrick McDaniel. IoT-Guard: Dynamic enforcement of security and safety policy in commodity IoT. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS)*, 2019.
- [12] Nicholas DeMarinis and Rodrigo Fonseca. Toward usable network traffic policies for IoT devices in consumer networks. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy (IoTS&P)*, pages 43–48, 2017.
- [13] Soteris Demetriou, Nan Zhang, Yeonjoon Lee, XiaoFeng Wang, Carl Gunter, Xiaoyong Zhou, and Michael Grace. Hanguard: Sdn-driven protection of smart home WiFi devices from malicious mobile apps. In *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, pages 122–133, 2017.
- [14] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring how privacy and security factor into IoT device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI)*, 2019.
- [15] Earlence Fernandes, Jaeyeon Jung, and Atul Prakash. Security analysis of emerging smart home applications. In *Proceedings of the 37th IEEE Symposium on Security and Privacy (SP)*, pages 636–654, 2016.
- [16] Earlence Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash. Flowfence: Practical data protection for emerging IoT application frameworks. In *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*, pages 531–548, 2016.
- [17] Earlence Fernandes, A. Rahmati, Jaeyeon Jung, and Atul Prakash. Security implications of permission models in smart-home application frameworks. *IEEE Security Privacy Magazine*, 15(2):24–30, 2017.
- [18] Earlence Fernandes, Amir Rahmati, Jaeyeon Jung, and Atul Prakash. Decentralized action integrity for trigger-action IoT platforms. In *Proceedings of the 25th Annual Network and Distributed System Security Symposium (NDSS)*, 2018.
- [19] Christine Geeng and Franziska Roesner. Who’s in control?: Interactions in multi-user smart homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI)*, pages 268:1–268:13, 2019.
- [20] Garrett M. Graff. Now a dorm room minecraft scam brought down the Internet. <https://www.wired.com/story/mirai-botnet-minecraft-scam-brought-down-the-internet/>, December 2017.
- [21] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlence Fernandes, and Blase Ur. Rethinking access control and authentication for

- the home internet of things (IoT). In *Proceedings of the 27th USENIX Security Symposium (USENIX Security)*, pages 255–272, 2018.
- [22] Kai-Hsiang Hsu, Yu-Hsi Chiang, and Hsu-Chun Hsiao. Safechain: Securing trigger-action programming from attack chains. *IEEE Transactions on Information Forensics and Security*, 14(10):2607–2622, 2019.
- [23] Yunhan Jack Jia, Qi Alfred Chen, Shiqi Wang, Amir Rahmati, Earlene Fernandes, Z. Morley Mao, and Atul Prakash. ContextIoT: Towards providing contextual integrity to applied IoT platforms. In *Proceedings of the 24th Annual Network and Distributed System Security Symposium (NDSS)*, 2017.
- [24] Thorin Klosowski. Automation showdown: IFTTT vs Zapier vs Microsoft Flow. <https://lifelifehacker.com/automation-showdown-ifttt-vs-zapier-vs-microsoft-flow-1782584748>, 2016.
- [25] Xianghang Mi, Feng Qian, Ying Zhang, and XiaoFeng Wang. An empirical characterization of IFTTT: Ecosystem, usage, and performance. In *Proceedings of the 2017 Internet Measurement Conference (IMC)*, pages 398–404, 2017.
- [26] B. Montagu, B. C. Pierce, and R. Pollack. A theory of information-flow labels. In *IEEE 26th Computer Security Foundations Symposium (CSF)*, 2013.
- [27] Andrew C. Myers and Barbara Liskov. Complete, safe information flow with decentralized labels. In *IEEE Symposium on Security and Privacy (Cat. No.98CB36186)*, 1998.
- [28] Chandrakana Nandi and Michael D. Ernst. Automatic trigger generation for rule-based smart homes. In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security (PLAS)*, pages 97–102, 2016.
- [29] Sukhvir Notra, Muhammad Siddiqi, Hassan Habibi Gharakheili, Vijay Sivaraman, and Roksana Boreli. An experimental study of security and privacy risks with emerging household appliances. In *Proceedings of the 2014 IEEE Conference on Communications and Network Security*, pages 79–84, 2014.
- [30] Danny Palmer. Mirai botnet adds three new attacks to target IoT devices. <https://www.zdnet.com/article/mirai-botnet-adds-three-new-attacks-to-target-iot-devices/>, May 2018.
- [31] Jonathon W. Penney. Whose speech is chilled by surveillance? <https://slate.com/technology/2017/07/women-young-people-experience-the-chilling-effects-of-surveillance-at-higher-rates.html>, July 2017.
- [32] Roei Schuster, Vitaly Shmatikov, and Eran Tromer. Situational access control in the Internet of Things. In *Proceedings of the 25th ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1056–1073, 2018.
- [33] Anna Kornfeld Simpson, Franziska Roesner, and Tadayoshi Kohno. Securing vulnerable home IoT devices with an in-hub security manager. In *Proceedings of the 15th IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 551–556, 2017.
- [34] Vijay Sivaraman, Hassan Habibi Gharakheili, Arun Vishwanath, Roksana Boreli, and Olivier Mehani. Network-level security and privacy control for smart-home IoT devices. In *Proceedings of the 11th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 163–167, 2015.
- [35] Milijana Surbatovich, Jassim Aljuraidan, Lujó Bauer, Anupam Das, and Limin Jia. Some recipes can do more than spoil your appetite: Analyzing the security and privacy risks of IFTTT recipes. In *Proceedings of the 26th International World Wide Web Conference (WWW)*, 2017.
- [36] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. “I don’t own the data”: End user perceptions of smart home device data practices and risks. In *Proceedings of the 15th Symposium on Usable Privacy and Security (SOUPS)*, 2019.
- [37] Yuan Tian, Nan Zhang, Yueh-Hsun Lin, XiaoFeng Wang, Blase Ur, Xianzheng Guo, and Patrick Tague. SmartAuth: User-centered authorization for the Internet of Things. In *Proceedings of the 26th USENIX Security Symposium (USENIX Security)*, pages 361–378, 2017.
- [38] Qi Wang, Wajih U. Hassan, Adam Bates, and Carl Gunter. Fear and logging in the internet of things. In *Proceedings of the 25th Annual Network and Distributed System Security Symposium (NDSS)*, 2018.
- [39] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata McDonough, and Yang Wang. Privacy perceptions and designs of bystanders in smart homes. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):59:1–59:24, November 2019.
- [40] Tianlong Yu, Vyas Sekar, Srinivasan Seshan, Yuvraj Agarwal, and Chenren Xu. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network

Rank	Channel (# of participants)
1	Notifications (21)
2	Email (18) Weather Underground (18)
4	Date and Time (16)
5	Google Sheets (15)
6	IFTTT (13) SmartThings (13)
8	Amazon Alexa (12) Location (12)
10	RSS Feed (11)

Table 7: The top ten channels in rank order by the number of participants using them.

Rank	Channel :: Trigger (# of participants)
1	Amazon Alexa :: Say a specific phrase (12) IFTTT :: New trigger or action published by service (12)
3	Date and Time :: Every day at __ (11) Weather Underground :: Tomorrow’s forecast calls for __ (11)
5	Button Widget :: Button press (10) Google Assistant :: Say a simple phrase (10) RSS Feed :: New feed item (10)
8	Location :: You exit an area (9)
9	Location :: You enter an area (8)
10	Date and Time :: Every day of the week at __ (7)

Table 8: The top ten triggers in rank order by the number of participants using them.

security for the Internet-of-Things. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks (HotNets)*, pages 5:1–5:7, 2015.

- [41] Eric Zeng and Franziska Roesner. Understanding and improving security and privacy in multi-user smart homes: A design exploration and in-home user study. In *Proceedings of the 28th USENIX Security Symposium (USENIX Security)*, pages 159–176, 2019.

A Appendix A: Frequently-Used Applet Components

Participants had a total of 743 IFTTT rules. Although there was substantial diversity in their rules, certain channels, triggers, actions, and trigger-action pairs were used by several participants. The top ten channels, triggers, actions, and trigger-action pairs (i.e., used by the largest number of participants) are shown in Tables 7 to 10.

B Appendix B: Descriptions of Information-Flow and Semantic Labels

Each trigger and action is labelled in three ways, based on its: integrity properties, secrecy properties, and semantic meaning.

Rank	Channel :: Action (# of participants)
1	Notifications :: Send a notification from the IFTTT app (2)
2	Email :: Send me an email (18)
3	Google Sheets :: Add row to spreadsheet (15) SmartThings :: Switch off (10)
5	Phone Call (US Only) :: Call my phone (9) SmartThings :: Switch on (9)
7	Google Calendar :: Quick add event (7) SMS :: Send me an SMS (7)
9	Android Device :: Set ringtone volume (6) Dropbox :: Add file from URL (6) Gmail :: Send me an email (6)

Table 9: The top eleven actions in rank order by the number of participants using them.

Rank	Trigger→Action Pair (# of participants)
1	IFTTT :: New trigger or action published by service → Email :: Send me an email (12)
2	Amazon Alexa :: Say a specific phrase → Phone Call (US Only) :: Call my phone (7) Weather Underground :: Tomorrow’s forecast calls for __ → Notifications :: Send a notification from the IFTTT app (7)
4	RSS :: New feed item → Email :: Send me an email (6)
5	Space :: ISS passes over a specific location → Notifications :: Send a notification from the IFTTT app (5) Weather Underground :: Today’s weather report → Notifications :: Send a notification from the IFTTT app (5)
7	Facebook :: You are tagged in a Facebook photo → Dropbox :: Add file from URL (4) Fitbit :: New sleep logged → Google Sheets :: Add row to spreadsheet (4) Google Assistant :: Say a simple phrase → SmartThings :: Switch off (4) Google Assistant :: Say a simple phrase → SmartThings :: Switch on (4) IFTTT :: New applet published by service → Email :: Send me an email (4)

Table 10: The top eleven trigger-action pairs in rank order by the number of participants using them.

Rank	Trigger Label → Action Label (# of participants)
1	News-ish → Log or notify (19)
2	Weather or time → Log or notify (15)
3	Voice command → Log or notify (14)
4	Sensing IoT device state → Log or notify (11) Environment sensing → Log or notify (11) Sensing online account state → Log or notify (11) Voice command → Change IoT device state (11)
8	Intentional trigger → Log or notify (8)
9	Actions with personal devices → Log or notify (7) Weather or time → Change IoT device state (7) Sensing online account state → Outgoing communication (7)

Table 11: The top eleven most prominent combinations of semantic trigger label – semantic action label pairs in rank order by the number of participants with this type of applet.

Although the labels themselves are descriptive, we include more thorough explanations of each label in Tables 12–14.

Information-Flow Label	Who could cause the trigger or action event to occur (<i>integrity</i>) or know that it occurred (<i>secrecy</i>)?
Public (pub)	<i>Secrecy</i> : Anyone in the world. E.g., an action that posts a public Tweet.
Untrusted (unt)	<i>Integrity</i> : Anyone in the world. E.g., the trigger “ <i>Android SMS :: Any new SMS received</i> ” (anyone can <i>send</i> the user an SMS).
Untrusted group (unt_g)	<i>Integrity</i> : Anyone in the world, but requiring group coordination. E.g., the trigger “ <i>reddit :: New top post in subreddit</i> ” could not be caused by a <i>single</i> untrusted user, but the users who could potentially coordinate to cause the event to occur are unrestricted.
Restricted Physical (rp)	<i>Secrecy</i> and <i>Integrity</i> : People with physical access to a particular space or device. E.g., an action that turns off a smart light has restricted physical secrecy and integrity, because anyone who can see the light could know that the action has occurred, and anyone with physical access to the device can control it.
Restricted Online (ro)	<i>Secrecy</i> and <i>Integrity</i> : A restricted group of online users. E.g., an action that creates a Facebook post visible to only friends has restricted online secrecy.
Trusted Physical (tp)	<i>Secrecy</i> and <i>Integrity</i> : Similar to restricted physical, but within a more trusted group. E.g., the trigger “ <i>Automatic Pro :: Entered an area</i> ” can be caused by someone with physical access to the car and the ability to drive it (typically only trusted friends or family).
Trusted Online (to)	<i>Secrecy</i> and <i>Integrity</i> : Similar to restricted online, but within a more trusted group. E.g., household members explicitly given access to remotely control or monitor a particular smart home device.
Trusted Other (t_oth)	<i>Integrity</i> : Trusted sources such as weather or news reports.
Trusted (t)	<i>Integrity</i> : The IFTTT user.
Private (priv)	<i>Secrecy</i> : The IFTTT user.

Table 12: Descriptions of information-flow labels. Shorthand for each semantic label, used in Appendix C, is shown in parentheses.

Semantic Trigger Label	Description	# Distinct Triggers
Weather or time (WT)	Official weather (e.g., not the weather sensed by a home IoT device) or a pre-specified, likely recurring time or date.	13
News-ish (N)	Actual news sources like the New York Times, updates from official sources about new products or apps, personal news such as shipment status updates or new RSS Feed items, or new posts on websites like Craigslist or Reddit.	35
Sensing IoT device state (DS)	Based on the state of an IoT device that is meant to be controlled by a person, such as whether a door is open, closed, locked, or unlocked.	32
Environment sensing (E)	Based on the state of an IoT device that is meant to reflect to the ground truth state of the environment it is in (including detecting motion, temperature, CO2, etc.).	19
Intentional trigger (I)	Pressing a particular button in a phone app or widget, or sending an SMS or Voicemail to IFTTT (which only has the purpose of acting as a trigger).	4
Voice command (V)	Triggers that can be caused by verbal interaction with a smart assistant. Includes items being added to Alexa lists and alarms going off, as well as arbitrary phrases spoken to the voice assistant.	10
Incoming communication (IC)	The IFTTT user receives communication from others. For example, receiving an SMS or an email.	6
Sensing online account state (OAcc)	Triggers that react to online account updates. Includes, for example, being tagged in a photo on Facebook, having new “liked” videos on YouTube, or a new sleep being logged by FitBit	24
Actions with personal devices (P)	Includes mainly location sensing (e.g., moving through space with a physical device), but also sensing that the user has sent an arbitrary SMS from their device, took a screenshot, etc. Includes location sensing for the IFTTT user’s device as well specified family members’ devices. Includes location inference via devices coming within range of a specific area (e.g., connecting to home WiFi).	15
Other automations (OAu)	“Stringify flow runs” or “Webhooks, receive a web request.”	2

Table 13: Descriptions of semantic labels of triggers. Shorthand for each semantic label, used in Appendix C, is shown in parentheses.

Semantic Action Label	Description	# Distinct Actions
Change IoT device state (DS)	Action alters or sets the state of an IoT device. E.g., turns lights on or off, locks or unlocks a door, turns on the thermostat.	59
Home security (DS:S)	A subset of “Change IoT device state” which affect IoT devices related to home security (e.g., door locks, surveillance cameras).	12
Lights (DS:L)	A subset of “Change IoT device state” which control lights.	20
Log or notify (L)	Creates a record of the trigger or notifies the user when the trigger occurs. Notifications can happen via notifications, emails, or phone calls where the action specifies that it is “to me.” Logs can be saved to cloud services such as Google Sheets or in other personal accounts like making a “Journal Entry” in Day One or adding a calendar event to a Google calendar.	19
Change personal device state (P)	Changes the state of a personal device (e.g., phone), for example by launching an app (maps, music), changing the phone volume, or turning the WiFi on or off.	8
Outgoing communication (OC)	Sends information to other people, including by sending email or SMS, or updating social media accounts.	10
Other automations (OAu)	These actions act as triggers for a Stringify flow, Webhooks request, Wink shortcut, or Nexia automation.	4

Table 14: Descriptions of semantic labels of actions. Shorthand for each semantic label, used in Appendix C, is shown in parentheses.

C Appendix C: Trigger and Action Labels

We reached consensus in our labelling of triggers and actions used in participants’ applets as described in Section 3. Tables 15 and 16 show all triggers and actions and their respective secrecy, integrity, and semantic labels.

Table 15: Semantic and information-flow labels of triggers used in participants’ applets.

Trigger channel :: trigger (# of rules with this trigger)	Semantic label	Secrecy label	Integrity label
Amazon Alexa :: Ask whats on your Shopping List (4)	V	rpUto	rp
Amazon Alexa :: Item added to your Shopping List (1)	V	rpUto	rpUto
Amazon Alexa :: Item added to your To Do List (4)	V	rpUto	rpUto
Amazon Alexa :: New song played (2)	V	rpUto	rpUto
Amazon Alexa :: Say a specific phrase (69)	V	rpUto	rp
Amazon Alexa :: Your Alarm goes off (7)	V	rpUto	rpUto
Amazon Alexa :: Your Timer goes off (2)	V	rpUto	rp
Ambient Weather :: Daily Rain rises above (1)	E	rp	rp
Android Device :: Connects or disconnects from any WiFi network (1)	P	rp	t
Android Device :: Connects to a specific WiFi network (1)	P	priv	t
Android Phone Call :: Any phone call answered (1)	IC	rp	t
Android SMS :: Any new SMS received (4)	IC	rp	unt
Android SMS :: Any new SMS sent (1)	P	rp	t
Android SMS :: New SMS received matches search (3)	IC	priv	unt
Apple App Store :: New app featured in a collection (1)	N	pub	t_oth
Apple App Store :: New app from search (1)	N	pub	unt t_oth
Apple App Store :: Top ten app goes on sale (1)	N	pub	unt_g
AppZapp :: Top App gone free in Google Play (2)	N	pub	unt_g
AppZapp :: Top App gone free in the Apple App Store (3)	N	pub	unt_g
Arlo :: Motion detected (1)	E	to	rp
August :: Lock locked (2)	DS	rpUto	rpUto
August :: Lock unlocked (1)	DS	rpUto	rpUto
Automatic Classic :: Check engine light turned on (3)	DS	rpUto	tp
Automatic Classic :: New trip completed (1)	DS	rpUto	tp
Automatic Pro :: Check engine light turned on (2)	DS	rpUto	tp
Automatic Pro :: Entered an area (1)	DS	rpUto	tp
Automatic Pro :: Exited an area (1)	DS	rpUto	tp
Automatic Pro :: Ignition turned off in area (4)	DS	rpUto	tp
Automatic Pro :: Ignition turned on (2)	DS	rpUto	tp
Automatic Pro :: Ignition turned on in area (2)	DS	rpUto	tp
Automatic Pro :: New trip completed (2)	DS	rpUto	tp
Best Buy :: New product in category (1)	N	pub	t_oth
Best Buy :: Product price changes (3)	N	pub	t_oth
Boxoh Package Tracking :: Any shipping status change (1)	N	priv	t_oth
Button widget :: Button press (29)	I	rp	t
Camera widget :: Any new photo (2)	P	rp	t
Classifieds :: New post from search (5)	N	pub	unt
Dash :: Check engine light turned on (3)	DS	rpUto	tp
Dash :: New trip completed (1)	OAcc	rpUto	tp
Date and Time :: Every day at (20)	WT	pub	t_oth
Date and Time :: Every day of the week at (10)	WT	pub	t_oth
Date and Time :: Every hour at (3)	WT	pub	t_oth
Date and Time :: Every month on the (2)	WT	pub	t_oth
Date and Time :: Every year on (1)	WT	pub	t_oth
Dominos :: Order out for delivery (2)	N	priv	t_oth
Dropbox :: New photo in your folder (2)	OAcc	ro priv pub	ro t unt
ecobee :: Thermostat enters Smart Home/Away (3)	DS	rpUto	rpUto
ecobee :: Thermostat indoor temperature is greater than (1)	E	rpUto	rpUto
ecobee :: Thermostat indoor temperature is less than (2)	E	rpUto	rpUto
ecobee :: Thermostat outdoor temperature is less than (1)	E	pub	rp
ESPN :: New game start (1)	N	pub	t_oth
ESPN :: New in-game update (2)	N	pub	t_oth
eWeLink Smart Home :: 1 Channel Plug turned on or off (1)	DS	rpUto	rpUto
eWeLink Smart Home :: 1 Channel Switch turned on or off (3)	DS	rpUto	rpUto
Facebook :: New photo post by you (1)	OAcc	priv pub ro	t
Facebook :: You are tagged in a photo (10)	OAcc	priv pub ro	unt
Facebook :: Your profile changes (2)	OAcc	priv pub ro	t
Fitbit :: Daily activity summary (3)	OAcc	pub	t_oth
Fitbit :: Daily goal not achieved by __:__(1)	OAcc	priv	t
Fitbit :: New sleep logged (8)	OAcc	priv	t
Fitbit :: New weight logged (1)	OAcc	priv	t
Gmail :: New email in inbox from (2)	IC	priv ro	unt
Gmail :: New email in inbox from search (5)	IC	priv ro	unt
Gmail :: New email in inbox labeled (4)	IC	priv ro	unt
Google Assistant :: Say a phrase with a number (1)	V	rpUto	rp
Google Assistant :: Say a phrase with a text ingredient (5)	V	rpUto	rp
Google Assistant :: Say a simple phrase (69)	V	rpUto	rp
Google Calendar :: Any event starts (2)	OAcc	ro to priv pub	t_oth t ro unt
Google Calendar :: Event from search starts (1)	OAcc	ro to priv pub	t_oth t ro unt

Google Calendar :: New event added (2)	OAcc	ro to priv pub	t_oth t ro unt
Google Wifi :: Device Connects (3)	P	tp	tp
Google Wifi :: Device Disconnects (5)	P	tp	tp
HomeSeer :: A device is turned off (1)	DS	rp to	rp to
HomeSeer :: A device is turned on (3)	DS	rp to	rp to
IFTTT :: Daily recommended Applet for you (3)	N	pub	t_oth
IFTTT :: New Applet published by service (9)	N	pub	t_oth
IFTTT :: New IFTTT update (6)	N	pub	t_oth
IFTTT :: New trigger or action published by service (23)	N	pub	t_oth
Instagram :: Any new photo by you (1)	OAcc	ro pub	t
iOS Calendar :: New event added to specific calendar (1)	OAcc	ro to priv pub	t_oth t ro unt
iOS Contacts :: Any new contact (9)	OAcc	priv	t
iOS Photos :: Any new photo (1)	P	rp to priv	t
iOS Photos :: New photo added to album (1)	OAcc	rp ro priv	t
iOS Photos :: New screenshot (2)	P	rp priv	t
Leeo :: Temperature below threshold (1)	E	rp to	rp to
Life360 :: First family member arrives at a specific place (2)	P	rp to	to
Life360 :: Last family member leaves a specific place (3)	P	rp to	to
Location :: You enter an area (12)	P	rp	t
Location :: You enter or exit an area (5)	P	rp	t
Location :: You exit an area (12)	P	rp	t
Manything :: Motion detected (2)	E	priv	rp
MyQ :: Door closed (1)	DS	rp	rp
MyQ :: Door opened (1)	DS	rp	rp
Nest Protect :: Battery is low (10)	DS	rp to	t_oth
Nest Protect :: Carbon monoxide emergency (3)	E	rp to	rp
Nest Protect :: Carbon monoxide warning (8)	E	rp to	rp
Nest Protect :: Smoke alarm warning (10)	E	rp to	rp
Nest Thermostat :: Nest set to Away (3)	DS	rp to	rp to
Nest Thermostat :: Nest set to Home (3)	DS	rp to	rp to
Nest Thermostat :: Temperature drops below (2)	E	rp to	rp to
Nest Thermostat :: Temperature rises above (2)	E	rp to	rp to
Netatmo Weather Station :: Temperature rises above (3)	E	pub	pub
NJ Transit :: New bus advisory (1)	N	pub	t_oth
NPR :: New story published (2)	N	pub	t_oth
OhmConnect :: An #OhmHour starts (2)	N	pub	t_oth
Phone Call (US only) :: Leave IFTTT any voicemail (3)	I	priv	t
ProPublica :: Congress is scheduled to vote on a bill (1)	N	pub	t_oth
ProPublica :: The president signs a new bill into law (5)	N	pub	t_oth
RainMachine :: Device is offline (1)	DS	rp to	rp to
reddit :: Any new post in subreddit (1)	N	pub ro	unt ro
reddit :: New hot post in subreddit (1)	N	pub ro	unt_g ro
reddit :: New post from search (1)	N	pub ro	unt ro
reddit :: New post saved by you (1)	OAcc	priv	t
reddit :: New top post in subreddit (1)	N	pub ro	unt_g ro
RSS Feed :: New feed item (14)	N	pub priv	unt t
RSS Feed :: New feed item matches (9)	N	pub priv	unt t
Slice :: Any new shipment (2)	N	priv	unt
Slice :: Shipment status changes (1)	N	priv rp	unt
SmartThings :: Any new motion (1)	E	to	rp
SmartThings :: Closed (7)	DS	rp to	rp
SmartThings :: Locked (3)	DS	rp to	rp
SmartThings :: Moisture detected (1)	E	rp to	rp
SmartThings :: Opened (14)	DS	rp to	rp
SmartThings :: Presence detected (4)	P	to	tp
SmartThings :: Presence no longer detected (2)	P	to	tp
SmartThings :: Switched off (7)	DS	rp to	rp
SmartThings :: Switched on (17)	DS	rp to	rp
SmartThings :: Temperature drops below (2)	E	rp to	rp
SmartThings :: Temperature rises above (2)	E	rp to	rp
SmartThings :: Unlocked (3)	DS	rp to	rp
SMS :: Send IFTTT an SMS tagged (3)	I	priv	t
SMS :: Send IFTTT any SMS (3)	I	priv	t
Space :: Astronomy picture of the day by NASA (1)	N	pub	t_oth
Space :: Breaking news by NASA (1)	N	pub	t_oth
Space :: ISS passes over a specific location (6)	N	pub	t_oth
Spotify :: New track added to a playlist (1)	OAcc	pub priv	t
Square :: Any new payment (1)	OAcc	rp	tp
Stringify :: Stringify Flow runs (7)	OAc	priv rp ro pub	t t_oth rp ro unt_g unt
The New York Times :: New article from search (1)	N	pub	t_oth
The New York Times :: New popular article in section (3)	N	pub	unt_g
TiVo :: SKIP segment detected (4)	DS	rp	rp
Todoist :: New completed task (2)	OAcc	priv	t
Tumblr :: Any new post (1)	OAcc	pub ro	t
Tumblr :: New like (1)	OAcc	priv ro pub	t
Twitter :: New liked tweet by you (1)	OAcc	pub ro	t
Twitter :: New tweet by a specific user (4)	N	pub ro	t_oth
Weather Underground :: Today's weather report (4)	WT	pub	t_oth
Weather Underground :: Tomorrow's forecast calls for (1)	WT	pub	t_oth
Weather Underground :: Tomorrow's weather report (1)	WT	pub	t_oth
Weather Underground :: Current condition changes to (4)	WT	pub	t_oth

Weather Underground :: Current pollen count rises above (10)	WT	pub	t_oth
Weather Underground :: Current UV index rises above (30)	WT	pub	t_oth
Weather Underground :: Sunrise (14)	WT	pub	t_oth
Weather Underground :: Sunset (1)	WT	pub	t_oth
Webhooks :: Receive a web request (23)	OAu	priv rp ro pub	t t_oth rp ro unt_g unt
WeMo Smart Plug :: Switched off (1)	DS	rp	rp
WeMo Smart Plug :: Switched on (1)	DS	rp	rp
Wireless Tag :: Temperature is too high (1)	E	rp	rp
Wireless Tag :: Temperature is too low (2)	E	rp	rp
YouTube :: New liked video (3)	OAcc	pub priv	t

Table 16: Semantic and information-flow labels of actions used in participants' applets.

Action channel :: action (# of rules with this action)	Semantic label	Secrecy label	Integrity label
abode :: Change mode (1)	DS:S	rpUto	rpUto
Android Device :: Launch Google Maps Navigation (2)	P	rp	t rp
Android Device :: Mute ringtone (4)	P	rp	rp
Android Device :: Play a specific song (3)	P	rp	t rp
Android Device :: Play music (1)	P	rp	t rp
Android Device :: Set ringtone volume (11)	P	rp	rp
Android Device :: Turn off WiFi (2)	P	rp	rp
Android Device :: Turn on WiFi (3)	P	rp	rp
Android Device :: Update device wallpaper (2)	P	rp	t
Android SMS :: Send an SMS (2)	OC	ro	t
Android Wear :: Send a notification (4)	L	rp	t unt
Arlo :: Start recording (2)	DS:S	to (toUrp)	to rp
Blink :: Arm system (4)	DS:S	rpUto	rpUto
Blink :: Disarm system (4)	DS:S	rpUto	rpUto
CNCT Life :: Toggle on/off (4)	DS	rpUto	rpUto
D-Link Smart Plug :: Turn off (2)	DS	rpUto	rpUto
D-Link Smart Plug :: Turn on (2)	DS	rpUto	rpUto
Day One :: Create Journal entry (5)	L	priv	t
Dropbox :: Add file from URL (9)	L	ro pub priv	ro t unt
Dropbox :: Create a text file (8)	L	ro pub priv	ro t unt
ecobee :: Resume thermostat program (2)	DS	rpUto	rpUto
ecobee :: Set thermostat comfort profile until next transition (4)	DS	rpUto	rpUto
Email Digest :: Add to daily email digest (4)	L	priv	t
Email Digest :: Add to weekly email digest (9)	L	priv	t
Email :: Send me an email (75)	L	priv	unt
Evernote :: Append to note (1)	L	priv to ro	t to ro
Evernote :: Create image note from URL (1)	L	priv to ro	t to ro
eWeLink Smart Home :: Turn 1 Channel Switch on or off (2)	DS	rpUto	rpUto
Facebook :: Upload a photo from URL (1)	OC	priv ro pub	t to ro
Fitbit :: Log your weight (1)	L	priv	t
Garageio :: Close garage door (7)	DS:S	rpUto	rpUto
Garageio :: Open garage door (5)	DS:S	rpUto	rpUto
Gmail :: Send an email (10)	OC	priv ro	unt t
Gmail :: Send yourself an email (1)	L	priv	unt t
Google Calendar :: Create a detailed event (1)	L	ro to priv pub	t t_oth ro to unt
Google Calendar :: Quick add event (16)	L	ro to priv pub	t t_oth ro to unt
Google Contacts :: Create new contact (2)	L	priv	t
Google Drive :: Upload file from URL (1)	L	priv to ro pub	t to ro unt
Google Photos :: Upload photo to album (2)	OC	priv to ro pub	t to ro unt
Google Sheets :: Add row to spreadsheet (50)	L	priv to ro pub	t to ro unt
Google Wifi :: Prioritize Device (3)	DS	to (toUrp)	to
Harmony :: End activity (6)	DS	rpUto	rpUto
Harmony :: Start activity (7)	DS	rpUto	rpUto
iOS Calendar :: Create a calendar event (1)	L	ro priv pub to	t ro to
iOS Health :: Log weight (1)	L	priv	t
iOS Reading List :: Add item to Reading List (2)	L	priv	t
iOS Reminders :: Add reminder to list (4)	L	priv ro to	t ro to
Leo :: Change light color (1)	DS:L	rpUto	rpUto
LIFX :: Blink lights (6)	DS:L	rpUto	rpUto
LIFX :: Breathe lights (1)	DS:L	rpUto	rpUto
LIFX :: Change color of lights (3)	DS:L	rpUto	rpUto
LIFX :: Toggle lights on/off (1)	DS:L	rpUto	rpUto
LIFX :: Turn lights off (5)	DS:L	rpUto	rpUto
LIFX :: Turn lights on (6)	DS:L	rpUto	rpUto
Lockitron :: Lock Lockitron (1)	DS:S	rpUto	rpUto
Lockitron :: Unlock Lockitron (2)	DS:S	rpUto	rpUto
Lutron Caseta and RA2 Select :: Activate scene (2)	DS:L	rpUto	rpUto
Lutron Caseta and RA2 Select :: Set light level (4)	DS:L	rpUto	rpUto
MagicHue :: Switch to dynamic mode for your Lights (2)	DS:L	rpUto	rpUto
Manything :: Start recording (2)	DS:S	rpUto	rpUto
Manything :: Stop recording (2)	DS:S	rpUto	rpUto
Nest Thermostat :: Set temperature (2)	DS	rpUto	rpUto
Nest Thermostat :: Set temperature range (1)	DS	rpUto	rpUto

Nest Thermostat :: Turn on fan for 15 minutes (1)	DS	rp	to	rp	to
Nexia :: Run a Nexia automation (2)	OAU	priv	rp	ro	pub
Noon Home :: Change scene (2)	DS:L	rp	to	rp	to
Noon Home :: Turn off room (2)	DS:L	rp	to	rp	to
Notifications :: Send a notification from the IFTTT app (117)	L	rp	to	t	
Notifications :: Send a rich notification from the IFTTT app (1)	L	rp	to	t	
Philips Hue :: Blink lights (3)	DS:L	rp	to	rp	to
Philips Hue :: Change color (2)	DS:L	rp	to	rp	to
Philips Hue :: Dim lights (8)	DS:L	rp	to	rp	to
Philips Hue :: Set a scene in a room (3)	DS:L	rp	to	to	
Philips Hue :: Toggle lights on/off (8)	DS:L	rp	to	rp	to
Philips Hue :: Turn off lights (3)	DS:L	rp	to	rp	to
Philips Hue :: Turn on color loop (1)	DS:L	rp	to	to	
Philips Hue :: Turn on lights (3)	DS:L	rp	to	rp	to
Phone Call (US only) :: Call my phone (19)	L	rp	to	unt	
Pinterest :: Add Pin to board (2)	OC	priv	ro	pub	t
Pocket :: Save for later (1)	L	priv	to	t	
Pushbullet :: Push a file (1)	L	priv	to	t	
Pushbullet :: Push a note (10)	L	priv	to	t	
RainMachine :: Start a program (2)	DS	rp	to	rp	to
reddit :: Submit a new text post (2)	OC	pub	ro	unt	ro
Slack :: Post to channel (14)	OC	ro	priv	ro	t
Smart Life :: Turn off (2)	DS	rp	to	rp	to
Smart Life :: Turn on (3)	DS	rp	to	rp	to
SmartThings :: Lock (4)	DS:S	rp	to	rp	to
SmartThings :: Switch off (20)	DS	rp	to	rp	to
SmartThings :: Switch on (33)	DS	rp	to	rp	to
SmartThings :: Unlock (6)	DS:S	rp	to	rp	to
SMS :: Send me an SMS (9)	L	priv	to	t	
Spotify :: Add track to a playlist (5)	OC	priv	pub	ro	t
Spotify :: Save a track (1)	L	priv	to	t	
Stringify :: Run a Stringify Flow (28)	OAU	priv	rp	ro	pub
TiVo :: Display message (7)	L	rp	to	to	
TiVo :: Send remote control key (7)	DS	rp	to	rp	
Todoist :: Create task (1)	L	priv	ro	t	ro
TP-Link Kasa :: Toggle (1)	DS	rp	to	rp	to
TP-Link Kasa :: Turn off (1)	DS	rp	to	rp	to
TP-Link Kasa :: Turn on (1)	DS	rp	to	rp	to
Twitter :: Post a tweet with image (1)	OC	ro	pub	t	
Twitter :: Update profile picture (2)	OC	pub	to	t	
VoIP Calls :: Call my device (5)	L	rp	to	t	
Webhooks :: Make a web request (15)	OAU	priv	rp	ro	pub
WeMo Insight Switch :: Toggle on/off (1)	DS	rp	to	rp	to
WeMo Smart Plug :: Turn off (7)	DS	rp	to	rp	to
WeMo Smart Plug :: Turn on (6)	DS	rp	to	rp	to
Wink: Nimbus :: Set dial label (26)	DS	rp	to	rp	to
Wink: Shortcuts :: Activate shortcut (7)	OAU	priv	rp	t	rp
Wyze :: Disable motion detection (1)	DS	to	to	to	
Wyze :: Enable motion detection (1)	DS	to	to	to	

D Appendix D: Full Survey Instrument

D.1 Survey Flow

- The survey flow included:
 - Informed consent procedures and instructions for downloading our browser extension, which collected information about participants' applets (not included in this document)
 - Several questions about participants' general use of IFTTT and preferences about their applets (General Questions)
 - Several sets of looping questions pertaining to specific applets (Looping Questions)
 - Questions pertaining specifically to secrecy and integrity (Explicitly Asking About Secrecy and Integrity)
 - Demographic and IUIPC scale questions (not included in this document)
- **Blue text** was not shown to participants.
- Answer choices are shown in *[italicized square brackets]* after each question.
- Questions included in the analysis in our paper are shown in **bold**, and the section number where these results are conveyed is included in parenthesis. Questions were omitted from discussion for various reasons, including: (1) looping questions for applet categories that only a few participants used (social media), (2) evidence that participants did not understand our questions in the way we intended them, (3) the questions were not relevant to our specific research questions, which shifted from our original vision based on our analysis of participants' applets.

D.2 General Questions

1. **(Section 4.5.1) How many of your IFTTT applets did you create yourself (as opposed to using ones others have created)? If you're not sure, please make your best guess.**
[Numeric dropdown]
2. **(Section 4.5.1) Do you prefer to create your own applet or search for one that already exists?**
[Prefer to create, Prefer to search, No preference]
3. Have you ever turned on (i.e., started using) an applet for one-time use or a specific event (e.g., to easily upload photos during a concert or trip)?
[Yes, No, Don't remember]
4. **(Section 4.5.1) How often do you turn on applets based on a friend's or colleague's recommendation?**
[Never, Rarely, Sometimes, Often, Don't remember]
5. Have you ever not turned on an applet (or turned off one that you were already using) because you thought it might be unsafe? (E.g., you were uncomfortable with the permissions it asked for, you thought it might compromise account security)
[Yes, No, Don't remember, I've never wanted to use an applet I thought was unsafe]
6. **(Section 4.5.3) (If yes to question 5) What were your concerns?**
[Free response]
7. Have you ever had an incident where an applet made you feel unsafe or you felt it violated your privacy (e.g., unlocked your door when you weren't home, posted a picture to Facebook that you didn't want there)?
[Yes, No, Not sure]
8. **(Section 4.5.3) (If yes to question 7) Please describe:**
[Free response]
9. Have you made any of the applets you've created publicly available?
[Yes, No, Don't remember, I haven't created any applets]
10. Do you use any other task automation services besides IFTTT? (For example, services like Tasker, Zapier, Stringify, Microsoft Flow, etc.)
[Yes, No, Don't remember]
11. **(If yes to question 10) Which services?**
[Free response]
12. **(If yes to question 10) Can you integrate IFTTT applets within the other service's programs? For example, the service Stringify allows the user to add IFTTT applets as part of a Stringify flow.**
[Yes, No, Don't know]
13. **(If yes to question 12) Do you use this feature with any of your applets?**
[Yes, No, Don't remember]
14. How much do you agree with the following statements?
[Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree]
 - (a) **(Section 4.5.2) The applets I turn on behave as I would expect from their description.**
 - (b) **(Section 4.5.4 and Table 6) I am comfortable telling my friends what applets I use.**
 - (c) **(Section 4.5.4 and Table 6) I am comfortable telling my colleagues what applets I use.**
 - (d) **(Section 4.5.4 and Table 6) I am comfortable with anyone knowing what applets I use.**
 - (e) I would be upset if an applet triggered when I didn't intend it to.

- (f) A stranger could trigger some of my applets.
 - (g) If an applet I was using didn't do what I thought it would do when I installed it, I would notice right away.
 - (h) If an applet weren't that useful, I would turn it off or delete it.
 - (i) **(Section 4.5.2) I think the applets that I have turned on are safe to use.**
 - (j) I have been concerned about the permissions an applet asked for.
 - (k) When choosing an applet to turn on, my primary criterion is how useful it will be.
15. (If “strongly agree” for statement 11 in question 14.) Since usefulness is not necessarily your primary criterion, what other considerations do you have?
[Free response]
16. Have you ever tried to link the behavior of multiple applets (“chain” them together)? For example, if you get close to your house, your thermostat is set to home mode, and if your thermostat is set to home mode, your lights will turn on.
[Yes, No, Don't remember]
17. (If yes to question 16) Which applets?
[Free response]
18. (If yes to question 16) Did it work as expected?
[Yes, Sometimes yes sometimes no, No, Don't remember]
19. (If “no” or “sometimes yes, sometimes no” to question 18) Please explain what went wrong:
[Free response]
20. Have you ever **unintentionally** made a “chain” between applets? For example, if the temperature gets above a threshold, set the thermostat to cool, and if the thermostat is set below a specific temperature, open the windows.
[Yes, No, Don't remember]
21. (If yes to question 20) Which applets?
[Free response]
22. Have you ever **manually** deleted anything that was posted **automatically** (e.g., a social media post or cloud storage update) by an applet?
[Yes, No, Don't remember, I don't use this type of applet]
23. **(Section 4.5.3) (If yes to question 22) Which applet created the post, and why did you want to delete it? If you do not remember all the exact details, please explain as much as you can.**
[Free response]

24. Would you consider some applets to be more sensitive than others (i.e., you care more about who knows about them or who can trigger them)?
[Yes, No, Not sure]

25. (If yes to question 24) Which applets?
[Free response]

D.3 Looping Questions

D.3.1 Looping Set 1: Applets Using Physical Devices

Since you had applets that used physical devices, we will now ask a set of questions about physical devices, for up to 5 devices.

(The next set of questions (26-29) loops up to five times, for a randomly-chosen set of the participant's applets that used physical devices, as determined by service categories on the IFTTT website)

Asking about physical device [#] of up to 5.

26. Consider the applet “[applet title]”. In which room is the [device] device used in the applet [trigger/action]? If the device has multiple components/sensors/etc. (e.g., a smart hub or lighting system), please list all the rooms that you can remember.
[Free response]
27. Does more than one person in the household commonly access the [device] device used in the applet [trigger/action] (either online or in person)? “Access” could mean performing an activity that the device senses; e.g., opening a door that has a sensor attached or walking into a room with a motion sensor.
[Yes, No, Not sure, I live alone]
28. Does the location of the [device] device make you more protective of who knows about the applet “[applet title]”? (For example, an applet that unlocks a ground floor window might be considered more sensitive than one that unlocks a second floor window.)
[Yes, No, Not sure]
29. (If yes to question 28) Please explain:
[Free response]

D.3.2 Looping Set 2: Cloud Storage Applets

Since you had applets that use cloud storage, we will now ask a set of questions about the details of the cloud storage, for up to 5 applets.

(The next set of questions (30-31) loops up to five times, for a randomly-chosen set of the participant's applets that used cloud storage, as determined by service categories on the IFTTT website)

Asking about cloud applet [#] of up to 5.

30. (Section 4.2) Is the file or folder that the applet “[applet title]” updates accessible only to you, or is it shared with others (e.g., housemates, family)?

[Only me, A group, Don’t remember]

31. How often do you check the file or folder used in this applet to see the updates?

[Never, Rarely, Sometimes, Often, Don’t remember]

D.3.3 Looping Set 3: Social Media Applets

You have some applets where both the trigger and action use social media or blogging services. We will ask a small set of more detailed questions for up to 5 applets.

(The next set of questions (32-33) loops up to 5 times, for a randomly-chosen set of the participant’s applets that used cloud storage, as determined by service categories on the IFTTT website)

Asking about social media applet [#] of up to 5.

32. In the applet “[applet title]”, do the people who follow your [service associated with action] account also follow your [service associated with trigger] account?

[Yes, No, Not sure]

33. (If yes to question 32) What are the main differences between the audiences, in your view?

[Free response]

D.3.4 Looping Set 4: Violating Applets

You will now be asked a set of questions about your thoughts and perceptions of various side-effects applets can have. This set will repeat for up to 5 different applets.

(The next set of questions (34-44) loops up to 5 times for a randomly-chosen set of the participant’s applets that violate security principles, determined by information-flow analysis)

Asking detailed questions for applet [#] of [#].

34. Consider the applet “[applet title].” How likely is it that this applet could do the following, in your opinion?

[Definitely impossible, Probably impossible, Probably possible, Definitely possible]

- (a) Be triggered by someone outside of your household?
- (b) Cause monetary loss? (e.g., by increasing your electric bill or using up data) OR (e.g., by increasing your electric bill or causing you to replace devices more frequently) (if applet uses a physical device)
- (c) Cause an undesired event if you forget that you have it turned on?
- (d) Spread sensitive information online?
- (e) Cause you embarrassment?

- (f) (Displayed only if applet uses physical device, determined by service categories on the IFTTT website) Damage the physical device that it uses?

- (g) Be used to undermine your home security?

35. Have you ever experienced any of the above consequences or other harmful side-effects when using this applet?

[Yes, No, Not sure]

36. (Section 4.5.3) (If yes to question 35) Please describe the incident as best you recall, including which applet(s) were involved and what side-effects occurred.

[Free response]

37. (Section 4.5.4 and Table 5) Would you be upset if the applet contributed to the following situations occurring:

[Very Upset, Slightly Upset, Not Upset, This type of harm is impossible for this applet]

- (a) Private information gets posted online unintentionally, possibly embarrassing you.
- (b) You no longer directly control what files are downloaded from email or social media, possibly spreading malware on your computer.
- (c) (Displayed only if applet uses physical device, determined by service categories on the IFTTT website) Your electronic device is used in way it wasn’t designed for (such as being toggled on/off very rapidly), possibly reducing its longevity or damaging it.
- (d) Data gets uploaded to your cloud storage more often than you thought, possibly causing you to run out of space.
- (e) You consume more resources (e.g., electricity, phone data, cloud storage space), possibly increasing your bills or otherwise causing you to spend more money.

38. Did you consider the possibility of some of the preceding consequences when deciding to turn on the applet “[applet title]”?

[Yes, No, Don’t remember]

39. (If yes to question 38) Which questions?

[Free response]

40. Would you be upset if a friend knew you had the applet “[applet title]”?

[Yes, No, A little upset, Not sure]

41. Would you be upset if a colleague knew you had this applet?

[Yes, No, A little upset, Not sure]

42. Would you be upset if a *stranger* knew you had this applet?
[Yes, No, A little upset, Not sure]
43. Who is meant to be able to *purposefully* trigger this applet?
[Myself; Trusted individuals, such as my spouse; A wider circle of known individuals, such as my Facebook friends or house guests, Unknown third parties, such as websites or strangers]
44. **(Section 4.5.4) Would you be upset if someone not in the intended group purposefully triggered this applet?**
[Yes, No, Not sure]

D.4 Explicitly Asking About Secrecy and Integrity

D.4.1 Secrecy

(These questions pertain to up to 9 applets with secrecy violations plus 1 safe one, determined by information-flow analysis, selected randomly from the participant's applets. Titles of applets were displayed to participants.)

45. Some applets indirectly pass information from a smaller, more restricted group to a larger, more open one. For instance, an applet that posts to twitter when the user reaches a Fitbit fitness goal is leaking information that only the user knew to all their twitter followers. Out of your applets below, do any of them allow information to leak from a smaller audience to a larger one (whether inadvertently or on purpose)?
[Yes, No, Not sure]
46. **(Section 4.5.2) Thinking about the possible data leakage, has your desire to keeping using any of these ap-**

plets changed?

[Yes, I am more cautious of some applets now; No, my desire to use these applets has not changed; I'm not sure.]

47. **(Section 4.5.3) (If yes to question 46) Which applets?**
[Free response]

D.4.2 Integrity

(These questions pertain to up to 9 applets with integrity violations plus 1 safe one, determined by information-flow analysis, selected randomly from the participant's applets. Titles of applets were displayed to participants.)

48. Some applets allow devices or services usually accessed by members of a smaller, trusted group to be indirectly controlled by members of a larger, less trusted group. For instance, an applet that adds a photo to a Google Drive folder whenever the user is tagged in a Facebook photo is essentially allowing any of the user's Facebook friends to add files to the user's private folder (which by default only the user would have access to). Out of your applets below, do any of them allow less trusted groups to control devices or services usually controlled by more trusted groups (whether inadvertently or on purpose)?
[Yes, No, Not sure]
49. **(Section 4.5.2) Thinking about the possible loss of control, has your desire to keeping using any of these applets changed?**
[Yes, I am more cautious of some applets now; No, my desire to use these applets has not changed; I'm not sure.]
50. **(Section 4.5.3) (If yes to question 49) Which applets?**
[Free response]