Do Users Really Know Alexa? Understanding Alexa Skill Security Indicators

Yangyong Zhang
yangyongzhang.io@gmail.com
SUCCESS Lab, Texas A&M University

Raj Vardhan
raj_vardhan@tamu.edu
SUCCESS Lab, Texas A&M University

Phakpoom Chinprutthiwong
phakpoom.c@sskru.ac.th
Sisaket Rajabhat University

Guofei Gu
guofei@cse.tamu.edu
SUCCESS Lab, Texas A&M University

ABSTRACT

Amazon Alexa's booming third-party skill market has grown from 160 to 100,000 skills within three years. In this work, we make the first effort in demystifying the Alexa skill permission system by studying its security indicators. Our user study results show that most of the surveyed Alexa users did not understand the security implications of interacting with third parties via Alexa's voice user interface (VUI). Despite the potential risks of undesired resource sharing, more than two-thirds of the surveyed Alexa users considered third-party skills safe because they think these skills are Alexa- or Amazon-owned applications. Together with other uncovered deficiencies of skill security indicator designs, our study indicates a pressing need for a paradigm shift in designing security indicators for VUI systems.

CCS CONCEPTS

• **Security and privacy → Usability in security and privacy**; • **Human-centered computing →** *Auditory feedback.*

KEYWORDS

Alexa, permission, security indicators, user study

ACM Reference Format:
Yangyong Zhang, Raj Vardhan, Phakpoom Chinprutthiwong, and Guofei Gu. 2023. Do Users Really Know Alexa? Understanding Alexa Skill Security Indicators. In *ACM ASIA Conference on Computer and Communications Security (ASIA CCS '23), July 10–14, 2023, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3579856.3595795

1 INTRODUCTION

Amazon Alexa supports a rapidly growing third-party developer community. There are more than 100,000 skills published in the Alexa store [24] with more than 12,376 different developers. However, third-party skills have been reportedly posing threats to user privacy and security. For example, SkillExplorer [44] reports that some third-party skills have been requesting users' private information and eavesdropping without strictly adhering to the developer and platform policies.

To remediate the problem of potentially invasive skills, Amazon Alexa uses various security indicators to alert users regarding possible risks of using third-party skills. In this work, we refer to these indicators as *skill security indicators* which are associated with three methods of user resource sharing: skill permission, account linking [11], and skill inputs/outputs [44] (or skill I/O). A natural question arises - *how effective are skill security indicators in helping users make security- and privacy-preserving decisions?*

Previous research [31, 41, 47, 55] has extensively studied the usability of mobile- and web-based permission systems and proposed appropriate mitigation. However, these results are not directly applicable to scrutinize Alexa's voice user interface (VUI) design. This is because Alexa's VUI design is fundamentally different from existing visual- or tactile-oriented user interfaces (e.g., a touch screen interface) in several ways. First, it adopts outsourced and cloud-based application processing to accommodate the excessive computational power need for natural language processing (NLP). While mobile applications are primarily locally-hosted and source code level vetting [13] is a common practice, Alexa skills are usually hosted in third-party remote servers [17, 34], which are neither monitored nor checked by the platform. This incurs code update vulnerabilities [34, 58] which allow arbitrarily content change and aggressive use of acquired user resources. Second, VUI is invisible and single-tasking [40]. This is because many users do not have the access to a screen when using Echo devices. The invisibility forces VUI systems to minimize complex human-computer interactions and prefer turn-based simple tasks such as question answering and device controlling. Given a short response time and constrained or no run-time indicators (e.g., no prompted window exists in VUI), users often do limited cognitive processing [40, 50]. This leads to potential unawareness of skills' invasive behaviors such as run-time information collection.

In this work, we design two user studies to explore the never-before-studied skill security indicators. First, we conduct a user survey with 124 valid Alexa users to quantitatively test the effectiveness of skill security indicators in warning users against potential risks. Second, we perform a qualitative skill experiment by recruiting 41 valid Alexa users to use Alexa skills and conduct an interactive interview. In this experiment, we not only validate the results in the user survey but also scrutinize why users behaved in such ways. While the user survey aimed to collect user data based

ASIA CCS '23, July 10–14, 2023, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-0098-9/23/07…$15.00
https://doi.org/10.1145/3579856.3595795

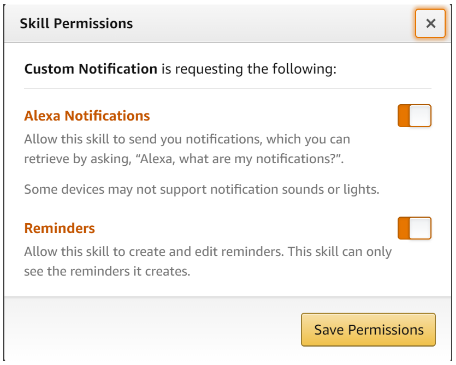
Skill Permissions

Custom Notification is requesting the following:

Alexa Notifications

Allow this skill to send you notifications, which you can retrieve by asking, "Alexa, what are my notifications?".

Some devices may not support notification sounds or lights.

Reminders

Allow this skill to create and edit reminders. This skill can only see the reminders it creates.

Figure 1: Permission prompt of a third-party skill named "Custom Notification" [12]. It is hard to differentiate two capabilities based on the descriptions. Note that changes have been made after we communicated with the Alexa skill team. Further information regarding these changes can be found in Section 6.1.

on their past experience, the interview-based skill experiment was designed to capture users' perception of skill security indicators in an immediate, in-depth manner.

Our study indicates dichotomous results. While many respondents understood the security implications of sharing conventional user resources (e.g., device address, phone number), most of them had little knowledge about risks incurred by VUI's unregulated back-end and voice-first features. For skill security indicators related to conventional capabilities (e.g., permissions on Android), more than half of the respondents were aware of the permissions. The comprehension rates (i.e., the percentage of respondents who comprehend the skill permission correctly) for these capabilities are high as well. However, we find that this knowledge is mostly inherited from users' previous experience in using mobile applications.

We further study the root cause of this seemingly harmless cognitive inertia because it is unclear why almost none of them was willing to take any actions. The result indicates that many Alexa users overlook and misunderstand the potential risks of using third-party skills. First, both the attention and comprehension rates for VUI-related security indicators are low. Second, most respondents did not understand who would be accessing their resources. For example, the user survey shows that most respondents (71%) thought it is Amazon (instead of the third-party skills) who requested skill permissions and account linking. Also, they thought Amazon and Alexa are trustworthy; hence the skills are safe to use. This explains why most respondents were not willing to take action. These findings illustrate that the skill security indicators require a significant redesign to improve their effectiveness. After analyzing these findings, we provide several short-term recommendations such as audio-based warnings and post-usage warnings[1], as well as discussion of long-term improvement strategies.

2 BACKGROUND AND RELATED WORK

2.1 Alexa Skills

2.1.1 Developing Third-party skills. In order to develop a third-party skill with the Alexa Skills Kit (ASK), two steps must be followed. The first step involves defining the language model for how the skill will expect users to interact with it, resulting in the production of an abstract parameter called intents through the use of unified speech processing. The second step is to build web services that can process the intents and other data received from the Alexa platform. These web services are hosted by the third parties themselves, and can be hosted using Amazon Web Service (AWS) accounts [17] or Alexa-hosted services. In either case, third-party developers have full control over their web services and can make updates without requiring re-certification [34].

2.1.2 Get Certification for Publishing a Third-party Skill. The certification process for skills developed for Amazon Alexa is designed to ensure that these skills meet certain security requirements [20]. For instance, the Alexa developer documentation specifies that skills should not include malicious hacking techniques such as phishing or Trojans. However, the way in which third-party skills behave after they have been published and how they manage acquired user resources is not subject to review or restriction by the Alexa store. This is because the current skill certification process does not require developers to provide the source code for their skills for review. Instead, third-party skills are hosted independently by the developers, and the content of these skills can be changed without any restrictions.

2.1.3 Skill Installation. Alexa users can install skills using two ways: (i) speak with skill-enabled devices and use voice commands to install a skill; (ii) browse and install skills through Alexa store (either web- or mobile-based). We find that many users prefer using the latter methods because the voice command based installation is limited in terms of the usability. First, users can only grant permissions or link their third-party accounts to skills via interfaces from web stores or mobile stores. Second, there are many skills with the same invocation names, and it is inaccurate to install a skill relying on voice commands alone [48].

Therefore, this study focus on users' installation experience through web- or mobile-based Alexa store. At the installation time, there are two ways for users to share access to their resources: skill permissions and account linking. Note that there is no installation time consent process for skill I/O. The skill permission is similar to mobile permission systems [41]; the key difference is that security enforcement is performed by the platform, which is located in the cloud. The account linking allows skill developers to access user-owned resources managed by third-parties, e.g., Facebook, Tesla, Google. The authorization is usually handled using OAuth 2.0 [11] or OpenID [51].

2.1.4 Interacting with Alexa. A user interacts with skills by speaking to Alexa-enabled devices such as Amazon Echo or smartphones with the Amazon Alexa application installed. In order to use a specific skill, users need to include the invocation name of a skill. For example, by speaking "Alexa, ask weather channel what's the weather.", a user will begin using the third-party skill named "The

[1] Most of the recommendations are provided in the project website, sites.google.com/view/dousersreallyknowalexa/.

Table 1: The descriptions provided in permission prompt (for skill permissions). The following descriptions are duplicates of the actual skill permission prompt messages.

Skill Permission	Description shown in the permission prompt
`Device Address`	Allow this skill to access the full postal address configured for your Alexa devices
`Country and Zipcode`	Allow this skill to access the country and postal code configured for your Alexa devices.
`Alexa Notification`	Allow this skill to send you notifications, which you can retrieve by asking, "Alexa, what are my notifications?"
`Read List`	Access to your Alexa lists
`Write List`	Permission to modify information on your Alexa lists
`Email`	Allow this skill to access the email address associated with your account
`Full Name`	Allow this skill to access the full name associated with your account
`Alexa Reminder`	Allow this skill to create and edit reminders. This skill can only see the reminders it creates
`First Name`	Allow this skill to access the first name associated with your account
`Amazon Pay`	Allow this skill to use Amazon Pay to make your payments. Amazon Pay will share your name, email and shipping address, but not your financial information, with the skill developer
`Mobile Number`	Allow this skill to access the mobile number associated with your account
`Location Service`	Allow this skill to access your location while the skill is in use

Weather Channel" and its invocation name is "weather channel". However, if the user speaks "Alexa, what is the weather?", the default weather service (provided by Amazon Alexa) is triggered. The default services usually already have access to user resources associated with their Alexa accounts. In this work, we focus on the third-party skills which need users' consent to access any user resources. After consenting to either skill permissions (e.g., the permissions listed in Table 1) or account linking, users will not be asked to grant any permissions when interacting with the skills. This ensures a smooth user experience.

2.2 Cognitive Inertia

Cognitive inertia is a psychological phenomenon that refers to the tendency for individuals to maintain their current beliefs, attitudes, and behaviors, even in the face of new information or experiences that may challenge or contradict those beliefs, attitudes, or behaviors [46]. This tendency is often referred to as "status quo bias," as it involves a preference for maintaining the current state of affairs [52].

Cognitive inertia can be influenced by various factors, such as past experiences, social influences, and cognitive biases [54]. For example, individuals may be more likely to exhibit cognitive inertia if they have previously experienced negative consequences from changing their beliefs or behaviors, or if they receive social reinforcement for maintaining their current beliefs or behaviors [36].

Cognitive inertia can have both positive and negative consequences. On the one hand, it may facilitate stability and consistency in an individual's beliefs, attitudes, and behaviors, which can be beneficial in certain contexts [46]. On the other hand, it may also lead to a resistance to change and a reluctance to consider new perspectives or alternatives, which can limit an individual's flexibility and adaptability [54].

2.3 Related Work

2.3.1 Voice Assistant Security. There are several research directions for studying Alexa security voice command attacks [28, 33, 57], voice squatting attacks [48, 58], and the Natural Language Understanding (NLU) related error-prone semantic interpretation [59]. However, there is little prior work studying permission (i.e., user resource management) issues in Amazon Alexa's backend application processing. Moreover, due to the usage of a new VUI design, security flaws are often caused by the gap between the user perception and the actual operations or design of Alexa. For example, voice squatting attacks exploit the gap between the users' intended voice commands and Alexa's ASR transcription output. Also, the masquerading attacks mentioned by Zhang et al. [58] is the result of the difference between user perception of when Alexa should stop and actual Alexa termination decision. Therefore, it is non-trivial to study the never-before-assessed Alexa skill permission system and its related skill security indicators.

Researchers have extensively studied various types of warnings and indicators across different platforms. To better understand how humans process warning messages, a commonly used model in the field is the Communication-Human Information Processing (C-HIP) model [56], which has been applied to examine different security indicators such as web [39] and mobile platforms [43]. The C-HIP model proposes that appropriate warning design can lead to quick user action. However, the effectiveness of security indicators in the popular Alexa skill platform, which employs unique Voice User Interface (VUI) features, remains an open question. In this study, we draw from previous research [29, 32, 43, 49] and apply a three-step C-HIP model in a user survey to scrutinize skill security indicators. Additionally, we aim to address the challenges posed by VUI, as it is difficult to provide the same level of feedback and guidance as traditional user interfaces. Consequently, we propose design recommendations for security indicators in the Alexa skill platform.

2.3.2 User Perception of Voice Assistants. Several studies have been done to learn about user perceptions on security and privacy aspects of voice assistants. Cho et al. [35] studies privacy and content customization of voice assistants and how users trust these systems. Abdi et al. [26] indicate that voice assistant users are not sure of how their data is being stored, processed, and shared. Huang et

Table 2: Security and privacy risks of Alexa and skill security indicator design.

	Security and Privacy Information	Skill Security Indicators		
		Skill Permission	Account Linking	Skill I/O
Conventional App Processing	R1. Requested scope and consequences	✓	✓	✗
	R2. Third-party identity	✓	✓	✗
VUI-related	R3. Dynamic Content	✗	✗	✓
	R4. Hidden behavior at back-end	✗	✗	✗
	R5. Run-time Information Collection	-	-	✗

✓: Warnings provided with either passively or proactively,
✗: No warnings provided in any form,
"-": Not applicable. This is because run-time information collection is not related to skill permissions and account linking.

al. [45] discusses the user perception of privacy risks in using shared smart speakers. In this work, starting from introducing skill security indicators, we conduct a more systematic approach to investigate the user perception of security and privacy risks when using voice assistants. Recent work [27, 38] further shows that voice assistants' privacy practices are worrisome.

3 SKILL SECURITY INDICATOR

In this section, to motivate the user studies of skill security indicators' effectiveness, we introduce the risks of using third-party skills with real-world examples. Then, we discuss about the detailed design of skill security indicators.

3.1 Risks incurred by Third-party Skills

After examining the user resources that Alexa skills can access, we identified five risks that can be classified into two categories: conventional capabilities commonly found in mobile apps, and voice user interface (VUI)-related risks that are new to users. This categorization provides a framework for understanding the potential security and privacy implications of using Alexa skills.

As shown in Table 2, the first category of risks includes two conventional capabilities that are commonly used in both mobile apps and Alexa skills. *Requested scope and consequences* (**R1**) refers to the access permissions requested by the app or skill and the potential consequences of granting such access. *Third-party identity* (**R2**) refers to the potential for third-party access to the user's identity information during the account linking process. To mitigate the risk of compromising their privacy and security, users should carefully review the permissions and account linking requests before granting access.

The second category includes three VUI-related risks that are critical to user security and privacy. These risks are *dynamic content* (**R3**), *hidden behavior at skill back-end* (**R4**), and *information collection* (**R5**), which occur during the distributed skill processing, and skill I/O (or run-time) respectively. These risks are unique to VUI, and they pose significant security and privacy challenges to users. In the following sections, we provide a detailed explanation of each of these three risks.

3.1.1 Distributed Skill Processing. To elaborate on R3 and R4 risks, we present a skill named MapNav [4] (published by ourselves) whose advertised functionality is to calculate the distance between one place to another. It requests the `Device Address` permission and the account linking to Google. Once a user linked her Google account, MapNav obtains the Google access token that can be used to access the user information as well as be used in other places.

R3: Dynamic Content. MapNav is able to play any speech feedback to users. We show an audio file-based approach in Listing 1. In Line 1 and 2, different audio URLs can be defined. Then, this skill can use any of them at Line 11 to play speech back to the user based on intent information at Line 5 and 6. This may incur unexpected and potentially inappropriate content such as information collection questions, abusive content, etc. Specifically, MapNav utilizes various predefined audio URLs, which can be chosen at runtime based on user intent. This may incur unexpected and potentially inappropriate content such as information collection questions, abusive words, etc.

```
1   const soundURL_1 = 'https://xxx.benign.com/file.mp3';
2   const soundURL_2 = 'https://xxx.malicious.com/file.
        mp3';
3   const StartSoundHandler = {
4       canHandle(handlerInput) {
5           return request.type === 'IntentRequest'
6           && request.intent.name === 'AnswerHandler';
                },
7       handle(handlerInput) {...
8         .addAudioPlayerPlayDirective
9       ('REPLACE_ALL', soundURL_1, expectedToken, 0,
             null)
10        .withSimpleCard('Example', 'Example')
11        .getResponse(); } };
```

Listing 1: Skill Source Code: Dynamic Content

R4: Hidden Behavior. This skill can also perform different hidden actions such as unauthorized resource sharing, gaining more resources than needed, etc. In detail, MapNav acquires the access tokens for Amazon and Google after interacting with the Alexa platform. Then it obtains the user data using these tokens. At the same time, the tokens can be sent out to other parties without any restriction. Also, this skill can not only acquire the location information on Google account profile but also unnecessary information such as a user's full name and email address. To do that, MapNav simply defines more fields in the request. We find that Google People API by default allows the requester to gain profile information that is not requested. Moreover, other conditional code executions such as behavior changes based on execution time are also feasible [34].

In Listing 2, we demonstrate how to share the gained resource and even gain more user data than needed. The skill gets the access tokens for Amazon (line 5) and Google (line 12) from the Alexa platform. Then it get the user data using the tokens in line 11 and line 21. At the same time, the tokens can be sent out to other parties without any restriction. Also, this skill can not only acquire the location information on Google account profile but also unnecessary information such as a user's true name and email address. For instance, line 2 shows how we can define more fields in the request. We find that Google People API by default allows us to gain profile information that is not requested. Moreover, other conditional code executions such as behavior changes based on execution time are also feasible [34].

```
1   const userFields = 'birthdays,addresses' ;
2   const APIkey = 'AIzaSyBaZjRAlgSb4B0FFYbQxxxxxxxx';
3   const permissions = ['read::alexa:device:all:address'
        ];
4   const consentToken = requestEnvelope.context.System.
        user.permissions
5   //get access token from Amazon for platform
        permission
6   && requestEnvelope.context.System.user.permissions.
        consentToken;
7   const deviceAddressServiceClient =
8   serviceClientFactory.getDeviceAddressServiceClient();
9   // get device address from Amazon
10  const address = await deviceAddressServiceClient.
        getFullAddress(deviceId);
11  var accessToken = System.user.accessToken;
12  //get access token for Google from Amazon Alexa
13  const urlUser =  'https://www.googleapis.com/oauth2/
        v1/userinfo?alt=json&access_token='+accessToken
        ;
14  // getting user's Google user ID
15  const userInfo = await (async () => { ...
16  return data;
17  const urlPeople = 'https://content-people.googleapis.
        com/v1/people/'+userOpenID+'?personFields='+
        userFields+'&key='+APIkey+'&access_token='+
        accessToken;
18  const people = await (async () => {
19  const fetchPeopleRes = await fetch(urlPeople);
20  const peopleJson = await fetchPeopleRes.json();
21  return peopleJson; //get data from Google
```

Listing 2: Skill Source Code: Access users' resource in their Google accounts

3.1.2 Covert VUI. Another risk is related to the questions raised by skills, e.g., a skill may ask questions like "What is your name?" or "What is your age?".

R5: Run-time Information Collection. The sensitive information collection at skill run-time can happen in any skills [34, 44]. Similar to the results reported by existing work [44], we find many skills do ask sensitive questions, and no skill security indicators are designed to warn users against such covert methods of information sharing. For example, when discussing this work with Amazon, we reported a skill named "Symptom Checker" (developed by "Infermedica") to have unjustified information collection. While this skill requested no permission or account linking, it would ask for unnecessary and unclaimed user information such as users' full names. Although this skill was soon removed from the Alexa store, it shows that there are still no effective protection mechanisms exist to protect users from the run-time information collection threat.

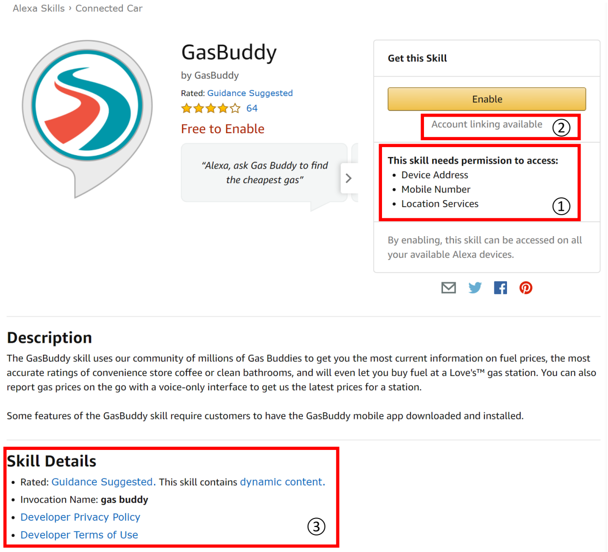

Alexa Skills › Connected Car

GasBuddy

by GasBuddy

Rated: Guidance Suggested

★★★★☆ 64

Free to Enable

"Alexa, ask Gas Buddy to find the cheapest gas"

Get this Skill

Enable

Account linking available

This skill needs permission to access:
- Device Address
- Mobile Number
- Location Services

By enabling, this skill can be accessed on all your available Alexa devices.

Description

The GasBuddy skill uses our community of millions of Gas Buddies to get you the most current information on fuel prices, the most accurate ratings of convenience store coffee or clean bathrooms, and will even let you buy fuel at a Love's™ gas station. You can also report gas prices on the go with a voice-only interface to get us the latest prices for a station.

Some features of the GasBuddy skill require customers to have the GasBuddy mobile app downloaded and installed.

Skill Details
- Rated: Guidance Suggested. This skill contains dynamic content.
- Invocation Name: **gas buddy**
- Developer Privacy Policy
- Developer Terms of Use

Figure 2: Skill homepage for "GasBuddy" skill. R1 (i.e., resource scope) is provided in ①, and R2 (i.e., third-party identity) is described with "by GasBuddy".

3.2 Indicator Design

Next, we describe the indicator design and what are the conveyed risks. We show that current skill security indicators are only designed to convey risks: R1, R2, and R3. In the following sections, we focus on providing an in-depth study for these three risk types, and we also discuss R4 and R5 in Section 6.

3.2.1 Skill Permissions. User resources associated with their Alexa accounts (e.g., device address, Amazon Pay) can be acquired via skill permissions. If one or more skill permissions are declared in a skill, two skill security indicators are shown to users. First, a *passive warning* (① in Figure 2) lists the skill permission names on the homepage. Second, a permission prompt window with all requested permissions (shown in Figure 1) is used to request users' consent. For each permission, a description is provided to help users understand the risks of consenting to the request. For example, the permission prompt in Figure 1 (in Section 1) explains what is `Alexa Notification` and `Alexa Reminder` [2]. Users can toggle the permission buttons to select what permissions to grant [22].

Similar to mobile permission systems [43], Alexa leverages skill permission prompts [22], together with the passive warning, to alert users regarding potential privacy- or security-invasive skills. Hence, skill permissions are designed to cover both R1 and R2.

3.2.2 Account Linking. User resources managed by non-Alexa entities can be accessed by account linking. For example, a skill can be linked to a user's Google account to access the associated user data, such as contact information. There are two skill security indicators designed for account linking. First, a passive warning (i.e., an

[2] A full list of skill permissions can be found in the project website [15]

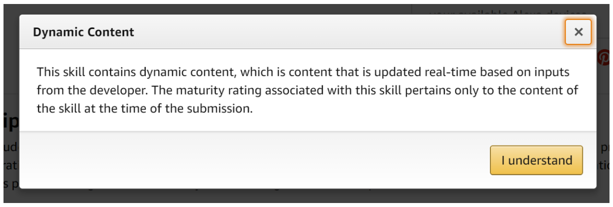

Dynamic Content

This skill contains dynamic content, which is content that is updated real-time based on inputs from the developer. The maturity rating associated with this skill pertains only to the content of the skill at the time of the submission.

Figure 3: Passive warning for skill I/O.

"account linking" icon shown as ② in Figure 2) will be shown on the skill homepage if the skill developer declares account linking. However, similar to the passive warning for skill permissions, no detailed information regarding the risks of using this function is displayed. Second, after the users clicked the "Enable" button, an account linking process will/can be triggered to redirect the users to an external log-in page. and users will be redirected to a log-in page for configuring the account linking. During this process, a log-in page warning (recommended by Amazon [11]) should exist to inform users of how the third-party skill access users' resources.

According to the Alexa developer documentation, " 'link accounts' means 'to get the user's permission to obtain ... the user data' " [21]. The underlying risk (i.e., R1 and R2) is that the third-party skill would gain access to users' sensitive resources associated with the linked account. To warn users against potentially invasive third-party skills, Amazon recommends that any implemented account linking process should let the users "view and accept any terms and conditions".

3.2.3 Skill I/O. Skills can also acquire user information through skill I/O [44]. When interacting with a user, skills can output a question and collect the information with the speech input. For example, a skill named Wiffy [10], which is designed to be a Wi-Fi password management skill, asks for users' phone numbers. Security indicators related to skill I/O are limited to a passive warning (i.e., the "Dynamic Content" shown in ③ in Figure 2). Specifically, in this warning (Figure 3), potential dynamic content is mentioned to warn users of the risks while interacting with third-party skills. This passive warning is clickable, but no prompt window or other proactive security indicator design is used. Also, this dynamic content warning exists in all skill homepages.

As shown in Figure 3, a skill's output can be dynamic (i.e., R3). Also, Amazon does not guarantee the effectiveness of the certification process provided by the Alexa market. It is admitted that a skill can change their behavior at any time. For example, a trivia skill [9] may change its output questions every time it is triggered. However, the change in skill output can be policy-violating or even malicious. For example, an originally normal question may contain inappropriate sexual comments [34] at a later time. Also, a neutral question could become a question aiming to collect private information.

3.3 Design Issues and Example Attacks

Skill Permissions. The first problem is the prompt window's permission-manager design [30]. This approach ensures a smooth user experience; however, it is less effective as a security indicator

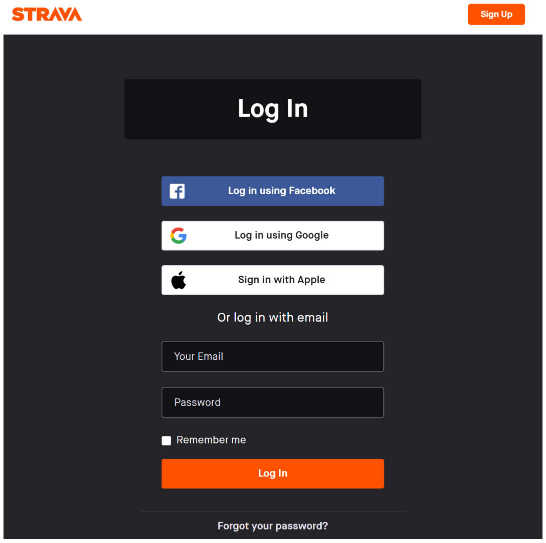

Figure 4: Account linking to Strava which is a popular exercise tracking application with more than 55 million users [25]. Four skills are found to have account linking to Strava (e.g., "Running history check for Strava" skill [23]). No scope declaration (including the requester information) or consent process is observed. After logged in, the account linking will automatically be done.

to convey R1 and R2 (i.e., requested scope and consequences & third-party identity). Users may not pay attention to these remindful skill permission prompts and take it for granted that these skill permission requests should be safe to consent. The second problem is the descriptions provided by permission prompts could be vague to users. As a result, a user might proceed with the skill without fully understanding the scope and security implications of granting skill permissions. For example, the description for `Alexa Reminder` capability [19] does not depict that a third-party skill can play audio messages automatically. This capability can potentially become invasive with the extensive control of when and how to play audio to users' private space[3]. Another example is that the Alexa list mentioned in `Read List` and `Write List` is unexplained. Users could underestimate the risks because skills can access both Shopping and To-Do Lists associated with users' Amazon accounts [18] using these two skill permissions.

Account Linking. The problem of existing account linking approach is that users may not check or understand the warning and then relate any risks to the skill. Skills' heterogeneous linking processes are implemented by different third-party authorization servers. This heterogeneity results in the question of whether log-in page warnings would always be effective in alerting users to R1 and R2. For example, as shown in Figure 4, no scope declaration or consent process is implemented by Strava. The missing log-in page warnings make it unclear what resources would be given away to the third-party skill. As a result, users' private information associated with their Strava account, such as home address, GPS location, daily activities, could be exposed unexpectedly. This is not

[3] The skill's backend code can be modified freely by the developers [34]

an isolated example because we find some other companies also did not have a well-implemented log-in page warning [3].

Skill I/O. The problem is that passive warnings, such as the click-to-open warning design used in this case, usually have very low attention rates [49, 55]. As a result, users may ignore this skill security indicator (shown in Figure 3) which is related to R3. Moreover, compared with visual interfaces, VUI is relatively new to users. Thus, if this indicator cannot convey all the risks properly (i.e., R3, R4, and R5 mentioned in Section 3.2), users may not notice any potential dangers when interacting with skills. For example, the Wiffy skill mentioned in Section 3.2.3 has a potentially policy-violating question to collect users' phone numbers. A user might not be aware of the potential problems of collecting their phone number (R5: run-time information collection) and potential unwanted information sharing (R4: hidden behavior at back-end). These problems lead to a potential ineffective skill security indicator design for skill I/O.

4 USER SURVEY

The user survey aims to quantitatively test the skill security indicators' effectiveness with the aforementioned risks considered. We start presenting the user survey by introducing the applied methodology and survey questions. Next, we describe user study details and the findings.

4.1 Methodology and Recruitment

When designing the user survey, we adhered to two design principles. First, following the best practice in studying security indicators [43], both user survey and skill experiment were used to correlate and justify the study results. Second, we leveraged C-HIP model [42, 56] as a guideline to design user studies for both studies. Specifically, the following three key steps in the C-HIP model are used:

- Attention: Do users pay attention to a skill security indicator? A user needs to switch focus from the primary task (i.e., installing a skill) to the security indicator, and she needs to focus on the security indicator for long enough to read and evaluate them.
- Comprehension: Do users understand the risk of granting permissions (for allowing resource access via skill permissions, account linking, or skill I/O) to third-party skills? Users need to understand the scope and implications of the permission.
- Behavior: Do skill security indicators influence users' installation decisions? Do users ever cancel installation because of the security indicators? Users should not install skills whose permissions exceed their comfort thresholds.

We did the recruitment and payment distribution of the user survey using Amazon Mechanical Turk (MTurk)[4] We used different MTurk recruitment filters to ensure the quality of the data collection. For example, we recruited MTurk workers who had high job approval rates (greater than 90%) and a sufficient number of approved jobs (greater than 500). Also, we recruited participants who are from the United States because our targeted Alexa language is "English (US)".

We initially recruited 150 respondents and paid $10 for each finished task. To ensure the quality of such online data collection, two attention check questions (e.g., "Can you open the survey link properly?") to help filter out bots and inattentive respondents. We also filtered out respondents who were inconsistent in answering the questions. For example, we ask a question at the beginning of the survey: if you own a smartphone? We do not include the answers from those who provided a negative response. We only recruited respondents who have used Alexa, and one needs a mobile app to initiate an Echo device or use the Alexa mobile application. Alexa users may use different portals. We asked the respondents to select their primary way of exploring and installing skills (either mobile, web, or voice command). We do not consider the results from those who only use voice commands to install skills. This was because skills with account linking or skill permissions cannot be installed via voice commands alone. As a result, 124 respondents' answers were collected. We assigned 18 minutes for a user to finish the survey. The average finish time was 14 minutes 17 seconds.

Among the 124 respondents, 60 were male, and 64 were female, with the remainder declining to identify their gender. Most of the respondents' age was between 29 and 39 (44%). Others' age distribution was: 23% between the ages of 18 and 28, 26% between the ages of 40 and 50, and 7% between the ages of 51 and 61. Moreover, 52% of the respondents use the web Alexa store as their primary Alexa portal. For education background, 69 of them own a bachelor's degree or higher.

4.2 Survey Questions

The survey questions are designed to focus on collecting users' perceptions of the skill security indicators based on their experience of using Alexa. In this study, we did not require the respondents to interact with Alexa. The user survey has 30 questions. The first part of it consists of 9 general questions asking about respondents' backgrounds, including their gender, education history, and familiarity with Alexa. The second part of the user survey consists of 21 questions regarding user experience related to skill security indicators. Six of the questions were randomly-picked skill permission quiz questions (see the full question list and results in Table 6 of Appendix C). These quiz questions were designed to study users' comprehension of skill permissions. To do that, we showed the respondents a skill permission prompt (with one permission request). Among the options, other than "I don't know" and "None of these", one or two correct options were provided. The correct options were reasonable inferences of the capabilities implied by the skill permissions.

4.3 Skill permissions

4.3.1 Attention. Do users pay attention to skill permission security indicators? Attention is the prerequisite of other steps in a C-HIP model. It is crucial to understand if Alexa users are paying attention to the skill indications, such as skill permissions in the first place. First, we asked the respondents if they ever used skills with

[4] This study's data collection was done between May 2020 to September 2020. All the studies are approved by the university Institutional Review Board (IRB). No personal identifiable information was collected or stored. To comply with the university COVID-19 guideline, no in-person user experiment was conducted.

Table 3: User survey results for skill permission attention rates.

Attention to Passive Warning		95% CI	Attention to Permission Prompt		95% CI
Looked at permission	39%	23% to 54%	Looked at permission	42.3%	20% to 59.2%
Did not look, but aware	51.7%	28% to 63%	Did not look, but aware	42%	24% to 61%
Checked Description	N/A	N/A	Checked Description	9.2%	7% to 15%
Was unaware of permission*	9.3%	4% to 13.5%	Was unaware of permission*	6.5%	3% o 12%

*: A respondent might forgot what they did in the past. We do not compare this result with skill experiment's result.
N/A: There was no description provided for the passive warnings shown in skill homepages.

permission(s). For respondents who answered yes, we then asked them[5], "For the last time when you installed an Alexa skill, which of the following properties of the skill were you aware of before you decided to enable it?" Next, for the respondents who chose the option of "skill permission", we further checked if they remember what they looked at? The answer could be either the passive warning or permission prompt. Note that we did not expect all the participants to remember what they looked at accurately. A more interactive and in-depth study will be applied in the skill experiment (Section 5). Lastly, if the respondents answered that they looked at the permission prompt, we further questioned them if they checked the permission descriptions (e.g., the descriptions in Figure 1).

We found that 13 respondents had never used skills with permissions or account linking. The reason could be that many skills published in the Alexa store only perform simple tasks like question answering. Thus, these skills usually do not request skill permissions or account linking. As shown in Table 3, even though a fair amount of Alexa users were aware of or looked at the permissions (either passive warning or permission prompt), they did not further check the detailed permission descriptions. The skill permission attention rates were low because only a few (9.2%) checked the details.

Finding 1: *Most respondents did not pay attention to skill permission details.* The results from the user survey show that some respondents did notice the skill permissions. However, only a small portion of them checked skill permission descriptions at least once when they were installing the skills. As suggested by the C-HIP model, this could prevent users from understanding the risks conveyed in the skill security indicators.

4.3.2 Comprehension. Do users understand the risks of granting skill permissions? It is essential to assess how Alexa users perceive the scope and implication of skill permissions. We asked each respondent 6 randomly-selected permission quiz questions. In each question, both the permission name and descriptions were displayed. The results indicate that most respondents understood conventional permissions such as `Device Address`. Also, we consider Alexa-specific skill permissions which are capabilities directly related to VUI functions such as Alexa list read/write, Alexa audio playing (e.g., `Alexa Reminders`). The results show that the comprehension rates for Alexa-specific permissions were low. For instance, `Alexa Reminders` has a comprehension rate of 31.8%. Also, users were confused at the difference between `Alexa Reminders` and `Alexa Notification`. Note that, the full list of permission descriptions and quiz results can be found in project website [15].

Also, for conventional permission quiz questions, there was no statistical correlation between age and the number of correct answers. However, for Alexa-specific permissions, there was a negative correlation ($r = -0.557, p < 0.001$); younger people were more likely to understand Alexa capabilities.

Finding 2: Compared to traditional permission models, a smaller percentage of Alexa users demonstrated a strong comprehension of Alexa-specific skill permissions, with only 44% providing correct responses overall. Users appear to be more familiar with traditional capabilities commonly used on other popular platforms, such as mobile phones. However, we did not assess the statistical significance of respondent groups that own or do not own mobile devices, as we excluded respondents who reported no use of mobile applications. Instead, we examined such behavior in the skill experiment conducted in Section 5.

Notably, 46.3% of respondents selected an incorrect option for the `Alexa Reminders` quiz question, which was the correct answer for `Alexa Notification`. While `Alexa Reminders` are usually user-set and intended to remind users of important tasks or events, `Alexa Notification` can be more intrusive since they are generated by third-party services and may be beyond the user's control.

We also asked respondents "Who requested and used the skill permission(s)?" to gauge their understanding of third-party involvement in the Alexa ecosystem. Surprisingly, a significant portion of Alexa users did not recognize the role of third parties in the ecosystem. In the user survey, 71% of respondents believed that either Amazon or Alexa was responsible for requesting skill permissions, while none of the respondents selected the correct answer, which was "the third-party skill".

Finding 3: *Many of the respondents thought that it is Amazon who requested the skill permission.* Some users did not consider the permissions carefully because they thought Amazon created the skills so that these skills can be trusted. This incurs a confused deputy problem. For example, if an evil skill can leverage users' misplaced trust to acquire private information or other resources. Specifically, because user-owned resources are often protected well by either Amazon or other platforms, it could be difficult for an attacker to gain these resources without becoming a skill developer. However,

[5] Some respondents may not understand these options. Thus, if they choose the option "I do not understand these options", an example homepage with these options marked was shown to them on the next page. They were given another chance to answer the question. The question can be found in Figure 5 of Appendix A.

once the users' trust in Amazon is falsely transferred to a skill controlled by a third party, the third-party could easily access users' resources.

4.4 Account Linking

4.4.1 Attention. Do users pay attention to the security indicators for account linking? There are two skill security indicators used in the account linking process: the passive warning before installing and the log-in page warning during the account linking process. We aim to examine if the respondents ever checked these two indicators. We first asked the respondents, "Have you ever used any skills with account linking?" We exclude respondents who never used skill account linking. As a result, 107 valid results were collected. Next, we raised the question, "Did you notice any warning messages when using the account linking?" We found that most of the respondents (81.3%) were aware of warnings regarding account linking in general. Fewer respondents (64.4%) specifically noticed the log-in page warning[6]. We used the skill experiment illustrated below to understand further why respondents paid more attention to account linking than skill permissions.

4.4.2 Comprehension. Do users understand the risks of the account linking? In the survey, for the 107 respondents who have used skills with account linking, we ask them regarding the passive warning of account linking, "what does 'account linking available' mean?" The result shows that 72% of the web-based Alexa respondents and 62% of the mobile-based respondents understood the conveyed risk of proceeding with the skill with account linking. We also showed the respondents three different log-in page warning examples (i.e., screenshots from Amazon, Google, and Best Buy account linking). There were 88.9% (32/36), and 72.7% (24/33) of the respondents who perceived the warning messages from Amazon and Google correctly. However, there were only 29% (11/38) respondents who checked Best Buy account linking warning understood it correctly (a detailed case study can be found in Appendix B). Another question we asked was about who requested account linking. There are 100 out of 107 respondents (93.5%) who answered either Amazon or Alexa. This further consolidates Finding 3.

Finding 4: *The respondents understood the log-in page warnings if it is well presented.* We found that most of the respondents understood the warnings provided in the account linking process. This indicates that, with proper indicator design, it is possible to inform the users regarding potential risks in a VUI setting.

4.5 Skill I/O

4.5.1 Attention and Comprehension. Do users pay attention to and understand the security indicator for skill I/O? As discussed in Section 3.2, only a passive warning is associated with skill I/O. Moreover, the passive warning was not shown to the user directly. Instead, a user has to click the "Dynamic Content" link in the Skill Detail section of the skill homepage (shown as ③ in Figure 2) to see skill I/O related warning message. In both experiments, we checked if any respondent ever clicked and checked the passive warning. The result shows that no respondents noticed or clicked it. We further want to check that assuming a respondent looked at the warning message, will she understand the risks conveyed? Thus, we presented them the"Dynamic Content" message and asked their understanding of that. As a result, 21 out of 124 respondents of the user survey and 4 out of 28 respondents from the skill experiment perceived it correctly.

Finding 5: *Alexa users did not check or understand skill I/O security indicators.* We found that the respondents are unfamiliar with skill I/O in terms of its potential risks or how it works. While VUI is new to any users, more effective and educational indicators should be provided to help users understand the risks in such a new user interface.

4.6 Behavior

We asked the respondents if they have ever decided not to install a skill? Only 6 respondents answered that they had declined the skill installation process. Four respondents claimed that they did not remember if they ever declined to install a skill. The others (114/124) provided negative answers. Next, for respondents who had ever decided not to install a skill, we asked them the reasons. The result shows that only 3 respondents were influenced by skill permission. No one declined to install or use a skill due to the security indicator of account linking. The potential reason could be that there are no actionable recommendations provided in the skill security indicator (i.e., the icon "Account Linking Available"). Thus, the respondents might not know what they are supposed to do even if they think the account linking can be dangerous. Also, no respondent reported that she was affected by the skill I/O security indicators.

5 SKILL EXPERIMENT

In the user survey, we found that risks regarding VUI are not well presented to users. To further verify these findings and explore users' mental models behind these findings, we initially recruited 45 respondents (4 excluded) to conduct a skill experiment. The participants in the two studies are recruited separately.

5.1 Experiment Design

Different from the questionnaire design of the user survey, the skill experiment asked respondents to install real-world Alexa skills and answer interview questions. There are two reasons for conducting the skill experiment. First, by asking users to use the skills, we can collect user data based on their fresh memory, which can be useful to confirm and explore the findings from the user survey. Second, the skill experiments were designed to gather nuanced data. With a longer experiment time assigned, we designed various open-ended questions to explore the mental models behind users' behaviors.

The first part includes general questions regarding education background, gender, etc. For the second part, we asked the respondents to use three different Alexa skills. After using each skill, they were interviewed regarding their experience in using the skills. The first skill, created by us, was used as a check-in skill to validate the respondents' ability to use Alexa and to help them warm up. The next two skills were selected from a pool of 9 skills covering various observed skill security indicator use cases. We assigned 30 respondents to each set of skill security indicators. We excluded 4 respondents who did not pass the check-in skill, resulting in 28

[6] We provided a note to help respondents understand what log-in page warnings are.

valid answers for each set of the indicators. Appendix A provides further details on the skills used.

5.2 Results

Skill Permission Attention. To validate Finding 1, we asked them the same questions as the user survey to check what they looked at when installing the skill. We collected 28 valid results from the skill experiment[7], and 17 (60.7%) of them claimed (at least for one time) that they noticed the skill permission. Five of them checked the detailed permission warnings. For those who did not check the detailed descriptions, we asked for the reasons. We found that many respondents tend to get rid of permission prompts as soon as possible. For example, one respondent's feedback is listed as follows,

> *I just keep clicking. The skill should be safe I think.(P-02, G2[8])*

The result consolidates Finding 1. The mental model of ignoring the skill permission prompts explains why many users were not checking skill permission details. However, it is still unclear about the reasons behind this mental model. Hence, we asked the respondents to answer a skill permission quiz with (similar to the user survey) 3 Alexa-specific and 3 conventional skill permissions. There are 21 (75%) respondents who answered all three conventional permission questions correctly. Among these 21 respondents, 9 of them reported that they never checked the skill permission details. We then asked them why they still knew the answer. We found that most of them (8/9) said the knowledge was inherited from their experience of using mobile applications. Thus, we conclude that there exists cognitive inertia that some Alexa users tend to only use what they know and resist changes.

Alexa-specific Permissions. Based on the quiz result, we found that only 1 (3.6%) of the respondents answered all Alexa-specific permission questions correctly. We also asked the respondents the difference between `Alexa Reminders` and `Alexa Notification`. After manually checking the answers, we found that most of the respondents (24/28) did not answer them correctly. Note there are 4 of these 24 respondents failed to provide a meaningful answer (blank or one-word answer). These results approve Finding 2.

Who uses my data? For respondents who used skill with skill permissions, we asked them who requested the skill permission; the result confirms Finding 3 that most users (20/28) did not realize that skills are not managed or owned by Amazon/Alexa. For respondents who experienced the account linking process, 9 out of 28 respondents (32.1%) answered that it was the third-party developer/skill who gained access to their resources. These results are consistent with Finding 3. For respondents who think it is Amazon or Alexa who requested and used the permission, we further asked them why they thought so. As a result, we found many of them thought that the Alexa and the skills are conceptually one entity. For example, one of the respondents answered:

> *I am talking to the Echo device and it is sold by Amazon. Also, I enabled the skill on Amazon website. Then it should be Amazon who did that. (P-033, G2)*

Log-in Page Warnings. We first asked the respondent if they noticed any security warnings during the account linking process. The result shows that 22 (out of 28) of the respondents were aware of the passive warning. Also, there were 17 respondents noticed the log-in page warning and 13 of these 17 respondents claimed they checked the details of the log-in page warnings. There were 2 out 10 respondents who tested "Running history check for Strava" reported that they were aware of log-in page warnings. We found that, compared with respondents who used the account linking process provided by Strava, the respondents who used skills with account linking to Amazon or Twitter performed significantly better in checking log-in page warnings (20% vs 83.3%; Mann-Whitney U Test, $U = 33, p < 0.05$). We checked the log-in page warning design of Strava and noticed that Twitter and Amazon's account linking processes are better presented with proper scope declaration and consent windows.

Skill I/O. After the respondents used skills without account linking and skill permissions, we asked them if they decided not to install the skill? As a result, only two out of 28 respondents have declined the skill installation. Both of them were assigned to test "Wiffy" skill. However, we found that they become cautious not because of the skill security indicators but other factors. For example, one respondent who declined to use Wiffy skill said that,

> *I don't feel comfortable giving out my Wi-Fi infomation to an app that I know little about. There isn't a concrete description of how the Wi-Fi password was stored, and if they can or cannot share it with others. (P-29, G2)*

We chose "Wiffy" based on the example policy-violating skill mentioned in SkillExplorer [44]. Guo et al. [44] showed that type of skill could be dangerous as it would collect users' private information covertly with skill I/O. The aforementioned result shows that the skill I/O security indicator failed to alert users regarding potential risks.

6 DISCUSSION

In this section, we will begin by outlining our suggested design recommendations for the security indicators of Alexa skills. Following that, we will explore various unresolved issues that have the potential to offer valuable insights for future research.

6.1 Design Recommendations

Skill Permission. The permission prompt should not be designed like a permission manager [30]. The current design implies a false impression that the skills have been installed and users are asked to remove any preset permissions if they want. We notice that users tend to skip the prompt window as soon as they can by clicking the "Save Permission" buttons. First, the skill permission should not be preset to be true. Second, the prompt window design should encourage the users to check the skill permission descriptions carefully. For example, similar to the mobile permission prompt design, the skill permissions' prompt window should ask users to choose either "Allow" or "Reject" for each permission. This can potentially increase the users' willingness of checking permission details.

Account Linking. Alexa relies on third parties to implement the log-in page warnings. However, we find that not all third-party

[7] The excluded results are from respondents who could not pass the first check-in skill.
[8] G1 is the participant group for user survey, and G2 is the participant group for skill experiment.

authorization servers provide well-defined and easy-to-read log-in page warnings. We first recommend that the Amazon Alexa store enforce a better log-in page design by including a stricter certification process for the account linking process. It should use a prompt window for each permission requested. Moreover, the scope should be clearly defined as well. Second, the passive warning of account linking should be more warning-like with detailed risk information included. Third, we recommend adding the scope of account linking to the skill permission prompt window. This is a feasible approach because the scope information can be acquired easily with the OAuth protocol. The benefit is that users can then check the warnings for account linking consistently rather than viewing them in different formats implemented by different third parties.

Skill I/O. It is not sufficient to only use a passive warning to alert the users to the risks of skill I/O. How a skill could leverage skill I/O should be clearly described in a more obvious manner. First, skills' I/O-related warning can be considered as a type of skill permission to be displayed in the permission prompt window. This would help increase the attention rate and help users understand what the risks could be. Another way to alert the users is to play audio warnings when users are interacting with the skills. Unfortunately, this is a challenging task because it will significantly decrease usability by reducing the effective interaction time. Audio-based warning messages could still be useful. For example, Alexa may use a different voice or accent when playing speech from a third-party skill. We also believe that post-usage warnings can be helpful. For example, in the current iOS permission system [16], a post-usage warning for background resource usage will warn users of how many times an app (such as Google Maps) has used the user's resource (e.g., location) in the background. Alexa may learn the lesson of leveraging post-usage warnings to help users understand implicit permission usage.

Alexa Skill Modifications. Shortly after completing the initial version of our study, we promptly communicated our findings and recommendations to the security team responsible for the Alexa skill platform. Recently, several changes have been observed, including the fact that skill permissions are no longer preset to be enabled. Instead, users are required to manually select the permissions they wish to grant. Additionally, the account linking log-in page has been enhanced to provide more detailed information about the user data that will be shared with third-party skills.

6.2 Open Problems

Cognitive inertia in VA. Users have been educated in using different mobile or web applications. Our study finds that many user behaviors are affected by the knowledge inherited from their past experience. For example, users tend to trust applications downloaded from Google Play App Store. However, there is a great difference between the traditional app store and the Alexa skill store. In specific, skills are not vetted based on their source code [34]. It could be dangerous to trust a skill even it is after the store vetting because the output from the skills can be malicious or policy-violating [34, 44]. How to better educate users and prevent harmful cognitive inertia remains an open question.

Third-party hosted skills. As many services are migrating to the cloud, it is becoming an increasingly challenging task to study security and privacy-related resource usage in modern applications. This is also a big problem for Alexa skills and other voice assistant platforms. The reason is that almost all these platforms let the third-party hosted skills on *any* cloud-based web services. As a result, it is difficult for the platforms to moderate the skill behaviors if they are not hosted in the same domain. Thus, the question is, how can a cloud-hosted skill be checked or monitored by the platforms or security researchers? This also leads to another question of how to enforce cloud-based permission systems?

6.3 Limitations, Ethics and Safety

We used a remote semi-moderated design in the skill experiment. In the future, to explore more about the usable skill security indicators, another way could be applying a moderated physical experiment to collect more insightful data [37, 53]. This could potentially help collect more details regarding users' thinking process. However, we argue that the methodology used in this work is also working well. By studying the methodologies from previous work [41, 58], we used various ways to ensure the user study is close to a real-world one. For example, we first recruited people who own voice assistant devices. Second, we used many screenshots and detailed scenario descriptions to mimic how users used Alexa skills in the real world. Our research was approved by the university IRB. Both the user studies and developed skills did not collect any personal identifiable information or any other sensitive user data. Moreover, we followed university guidelines over campus safety during the COVID-19 period, and no in-person experiment is conducted.

7 CONCLUSION

In this paper, we studied the security indicators of the popular Amazon Alexa platform. To understand how users perceive the skill security indicators, we performed two user studies: a user survey and a skill experiment. Our findings show that Alexa users pay little to no attention to some critical skill security indicators and often misinterpret the risks associated with skill permissions. Additionally, many users have misplaced trust in third-party skills wherein they may expect all skills to be safe because they are backed or owned by Amazon/Alexa. As this can lead to users falling victim to undesirable or even malicious skills, we discussed recommended changes to skill security indicators to improve their effectiveness. We hope that our findings can stimulate future research efforts in this emerging direction.

ACKNOWLEDGMENTS

This material is based upon work supported in part by NSF under Grant No. 1700544 and 1816497. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF.

REFERENCES

[1] [n. d.]. Alexa Skill: Angel Investor. https://www.amazon.com/Stoked-Skills-LLC-Angel-Investor/dp/B07GBLHKVR/.
[2] [n. d.]. Alexa Skill: Calm My Cat. https://www.amazon.com/Voice-Games-Relax-My-Cat/dp/B07JZ775V3/.

[3] [n. d.]. Alexa Skill: debugMining. https://www.amazon.com/%E8%B0%A2%E5%8B%87-debugMining/dp/B07FFKLPY2.

[4] [n. d.]. Alexa Skill: Mapnav. https://www.amazon.com/none-mapNav/dp/B0849QD5F4/.

[5] [n. d.]. Alexa Skill: Mediktor. https://www.amazon.com/Teckel-Solutions-S-L-Mediktor/dp/B075XRN21X/.

[6] [n. d.]. Alexa Skill: Shower Timer. https://www.amazon.com/CMM-Tech-Shower-Timer/dp/B07HWRBL8M.

[7] [n. d.]. Alexa Skill: The Magic Door. https://www.amazon.com/The-Magic-Door-LLC/dp/B01BMUU6JQ.

[8] [n. d.]. Alexa Skill: The Soap Box. https://www.amazon.com/Blutag-Inc-The-Soap-Box/dp/B07L129Y27/.

[9] [n. d.]. Alexa Skill: Twenty Questions . https://www.amazon.com/Amazon-Twenty-Questions/dp/B01C3CO48G.

[10] [n. d.]. Alexa Skill: Wiffy. https://www.amazon.com/hartman-Wiffy/dp/B06WVB8WM7/.

[11] [n. d.]. Amazon Alexa Account Linking Documentation. https://developer.amazon.com/en-US/docs/alexa/account-linking/account-linking-concepts.html.

[12] [n. d.]. Custom Notification Skill. https://www.amazon.com/Just-Reed-LLC-Custom-Notification/dp/B08G19NYMH/.

[13] [n. d.]. Google Play Protect. https://www.android.com/play-protect/,.

[14] [n. d.]. Proactive Events API . https://developer.amazon.com/en-US/docs/alexa/smapi/proactive-events-api.html,.

[15] [n. d.]. Project Website: Do Users Really Know Alexa? https://sites.google.com/view/dousersreallyknowalexa/. Accessed October 21, 2020.

[16] [n. d.]. What to do when your iPhone says an app has been using your location in the background. https://qz.com/1713581/what-to-do-when-an-iphone-app-is-using-\your-location-in-the-\background/.

[17] [n. d.]. Where to host an Alexa Skill. https://developer.amazon.com/en-US/docs/alexa/custom-skills/understanding-custom-skills.html. Accessed April 30, 2020.

[18] 2020. Alexa List in Alexa Skill Kit. https://developer.amazon.com/en-US/docs/alexa/custom-skills/access-the-alexa-shopping-and-to-do-lists.html.

[19] 2020. Alexa Reminder in Alexa Skill Kit. https://developer.amazon.com/en-US/docs/alexa/smapi/alexa-reminders-overview.html.

[20] 2021. Alexa Certification's Security Requirements . https://developer.amazon.com/en-US/docs/alexa/custom-skills/security-testing-for-an-alexa-skill.html.

[21] 2021. Alexa Developer Documentation: Account Linking Concept. https://developer.amazon.com/en-US/docs/alexa/account-linking/account-linking-concepts.html.

[22] 2021. Alexa Permission in Custom Skill. https://developer.amazon.com/en-US/docs/alexa/custom-skills/configure-permissions-for-customer-\information-in-your-skill.html.

[23] 2021. Alexa Skill: Running history check for Strava (unofficial). https://www.amazon.com/Running-history-check-Strava-unofficial/dp/B06XG4Z1N6.

[24] 2021. Voicebot.AI: Amazon Alexa Has 100k Skills. https://voicebot.ai/2019/10/01/amazon-alexa-has-100k-skills-but-momentum\-slows-globally-here-is-the-breakdown-by-\country/

[25] 2023. Strava Revenue and Usage Statistics . https://www.businessofapps.com/data/strava-statistics.

[26] Noura Abdi, Kopo M Ramokapane, and Jose M Such. 2019. More than smart speakers: security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*.

[27] Noura Abdi, Xiao Zhan, Kopo M Ramokapane, and Jose Such. 2021. Privacy norms for smart home personal assistants. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.

[28] Hadi Abdullah, Washington Garcia, Christian Peeters, Patrick Traynor, Kevin RB Butler, and Joseph Wilson. 2019. Practical Hidden Voice Attacks against Speech and Speaker Recognition Systems. *arXiv preprint arXiv:1904.05734* (2019).

[29] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *22nd {USENIX} Security Symposium ({USENIX} Security 13)*. 257–272.

[30] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 787–796.

[31] David Barrera, H Güneş Kayacik, Paul C Van Oorschot, and Anil Somayaji. 2010. A methodology for empirical analysis of permission-based security models and its application to android. In *Proceedings of the 17th ACM conference on Computer and communications security*. 73–84.

[32] Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter, and Manya Sleeper. 2014. Harder to ignore? Revisiting pop-up fatigue and approaches to prevent it. In *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*. 105–111.

[33] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. 2016. Hidden voice commands. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 513–530.

[34] Long Cheng, Christin Wilson, Song Liao, Jeffrey Young, Daniel Dong, and Hongxin Hu. 2020. Dangerous Skills Got Certified: Measuring the Trustworthiness of Skill Certification in Voice Personal Assistant Platforms. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 1699–1716.

[35] Eugene Cho, S Shyam Sundar, Saeed Abdullah, and Nasim Motalebi. 2020. Will deleting history make alexa more trustworthy? effects of privacy and content customization on user experience of smart speakers. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[36] Robert B Cialdini and Lloyd James. 2009. *Influence: Science and practice*. Vol. 4. Pearson education Boston.

[37] Neil S Coulson, Richard Smedley, Sophie Bostock, Simon D Kyle, Rosie Gollancz, Annemarie I Luik, Peter Hames, and Colin A Espie. 2016. The pros and cons of getting engaged in an online social community embedded within digital cognitive behavioral therapy for insomnia: survey among users. *Journal of medical Internet research* 18, 4 (2016), e5654.

[38] Jide Edu, Xavier Ferrer Aran, Jose Such, and Guillermo Suarez-Tangil. 2022. Measuring Alexa Skill Privacy Practices across Three Years. In *The Web Conference 2022*. ACM.

[39] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1065–1074.

[40] Kenneth Michael Farley, John O'Reilly, Leon Squire, and Rick Beasley. 2001. *Voice application development with VoiceXML*. Pearson Education.

[41] Adrienne Porter Felt, Erika Chin, Steve Hanna, Dawn Song, and David Wagner. 2011. Android permissions demystified. In *Proceedings of the 18th ACM conference on Computer and communications security*. 627–638.

[42] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*. 1–14.

[43] Adrienne Porter Felt, Helen J Wang, Alexander Moshchuk, Steve Hanna, and Erika Chin. 2011. Permission Re-Delegation: Attacks and Defenses.. In *USENIX Security Symposium*, Vol. 30. 88.

[44] Zhixiu Guo, Zijin Lin, Pan Li, and Kai Chen. 2020. SkillExplorer: Understanding the Behavior of Skills in Large Scale. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*. 2649–2666.

[45] Yue Huang, Borke Obada-Obieh, and Konstantin Beznosov. 2020. Amazon vs. my brother: How users of shared smart speakers perceive and cope with privacy risks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[46] Daniel Kahneman and Amos Tversky. 1979. On the interpretation of intuitive probability: A reply to Jonathan Cohen. (1979).

[47] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. 2012. A conundrum of permissions: installing applications on an android smartphone. In *International conference on financial cryptography and data security*. Springer, 68–79.

[48] Deepak Kumar, Riccardo Paccagnella, Paul Murley, Eric Hennenfent, Joshua Mason, Adam Bates, and Michael Bailey. 2018. Skill squatting attacks on amazon alexa. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 33–47.

[49] Kenneth R Laughery and Michael S Wogalter. 2014. A three-stage model summarizes product warning and environmental sign research. *Safety science* 61 (2014), 3–10.

[50] Christine Murad, Cosmin Munteanu, Benjamin R Cowan, and Leigh Clark. 2019. Revolution or evolution? Speech interaction and HCI design guidelines. *IEEE Pervasive Computing* 18, 2 (2019), 33–45.

[51] David Recordon and Drummond Reed. 2006. OpenID 2.0: a platform for user-centric identity management. In *Proceedings of the second ACM workshop on Digital identity management*. ACM, 11–16.

[52] William Samuelson and Richard Zeckhauser. 1988. Status quo bias in decision making. *Journal of risk and uncertainty* 1, 1 (1988), 7–59.

[53] William C Schmidt. 1997. World-Wide Web survey research: Benefits, potential problems, and solutions. *Behavior research methods, instruments, & computers* 29, 2 (1997), 274–279.

[54] Fritz Strack and Roland Deutsch. 2004. Reflective and impulsive determinants of social behavior. *Personality and social psychology review* 8, 3 (2004), 220–247.

[55] Christopher Thompson, Martin Shelton, Emily Stark, Maximilian Walker, Emily Schechter, and Adrienne Porter Felt. 2019. The web's identity crisis: understanding the effectiveness of website identity indicators. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 1715–1732.

[56] Michael S Wogalter. 2006. Communication-human information processing (C-HIP) model. *Handbook of warnings* (2006), 51–61.

[57] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 103–117.

[58] Nan Zhang, Xianghang Mi, Xuan Feng, XiaoFeng Wang, Yuan Tian, and Feng Qian. 2019. Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems. In *IEEE Symposium on Security and Privacy (SP)*. IEEE.

[59] Yangyong Zhang, Lei Xu, Abner Mendoza, Guangliang Yang, Phakpoom Chinprutthiwong, and Guofei Gu. 2019. Life after Speech Recognition: Fuzzing Semantic Misinterpretation for Voice Assistant Applications.. In *NDSS*.

A STUDY MATERIALS

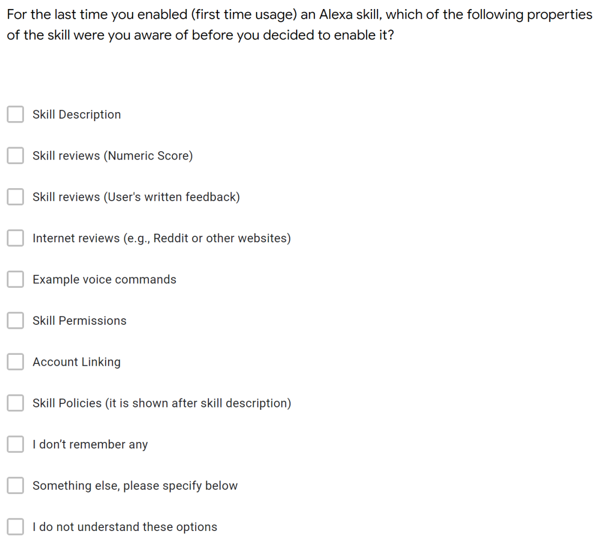

For the last time you enabled (first time usage) an Alexa skill, which of the following properties of the skill were you aware of before you decided to enable it?

□ Skill Description

□ Skill reviews (Numeric Score)

□ Skill reviews (User's written feedback)

□ Internet reviews (e.g., Reddit or other websites)

□ Example voice commands

□ Skill Permissions

□ Account Linking

□ Skill Policies (it is shown after skill description)

□ I don't remember any

□ Something else, please specify below

□ I do not understand these options

Figure 5: Survey question for skill permission attention.

The skills used for skill I/O are: Wiffy [10], Twenty Questions [9], and Calm My Cat [2].

Table 4: Skills used in the skill experiment of skill permission.

Skill Name	Requested Skill Permission
Mediktor [5]	Device Country and Postal Code
The Magic Door [7]	Alexa Notifications
The Soap Box [8]	Device Address, Full Name
	Email Address, Mobile Number

Table 5: Skills used in the skill experiment of account linking.

Skill Name	Account Linking
Running history check for Strava [23]	Strava
Shower Timer [6]	Instagram
Angel Investor [1]	Amazon

B CASE STUDY: BEST BUY SKILL

Alexa relies on third-party authorization servers to implement the log-in page warning. However, we have found that not all third-party servers provide well-defined and readable log-in page warnings. For instance, as shown in Figure 6, the Best Buy log-in page presents two problems. Firstly, the warning message on this page, which reads as follows,

"By linking your Best Buy account to Alexa, you will enable voice purchases from Best Buy without further action. To disable this function, you must delete the Best Buy skill."

is neither well-defined nor easily readable. The scope and implications of the warning are unclear, and the font size is small. Furthermore, the message is embedded in the log-in window, which makes it challenging for an Alexa user to either notice or comprehend the log-in page warning.

Secondly, we found that the Best Buy account linking can be done without the user's consent. We also identified a potential automatic log-in scenario. If a user previously logged into their Best Buy account and selected the "Keep me signed in" option, they will be automatically linked to a skill with an account linking request to Best Buy.

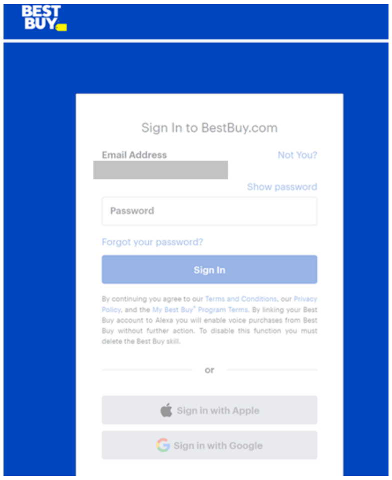

Sign In to BestBuy.com

Email Address
[redacted]

Password

By continuing you agree to our Terms and Conditions, our Privacy Policy, and the My Best Buy® Program Terms. By linking your Best Buy account to Alexa you will enable voice purchases from Best Buy without further action. To disable this function you must delete the Best Buy skill.

Figure 6: Example of automatic login during account linking.

C SKILL PERMISSION COMPREHENSION STUDY

In Table 6, we present the quiz questions and results for skill permission comprehension. We would like to note that during our communication with the Alexa team, we learned that the `Alexa Notification` permission was deprecated in late 2020, and a new version called the Proactive Events API [14] was introduced.

Table 6: User comprehension quiz questions for skill permissions.

Question	n	Options	Response
Device Address	69	✗ Amazon will send me advertisements to my physical address.	22 (31.9%)
		✓ The skill will know my full address associated with the device.	41 (59.4%)
		✗ The skill will know how I go to work.	4 (5.8%)
		✗ None of these.	0 (0%)
		I don't know.	2 (2.9%)
Country and Zipcode	68	✗ The skill can locate my phone location.	13 (19%)
		✓ The skill will know my zipcode associated with the device.	52 (76.4%)
		✗ My home address.	2 (3%)
		✗ None of these.	1 (1.5%)
		I don't know.	0 (0%)
First Name	69	✗ The skill will ask me whether I want to play a voice message.	2 (2.9%)
		✓ The skill can read the first name configured for my Alexa account.	48 (69.5%)
		✗ The skill may call me.	16 (23.2%)
		✗ None of these.	3 (4.3%)
		I don't know.	0 (0%)
Full Name	68	✗ Amazon can make purchases using my name.	1 (1.5%)
		✗ Amazon will know the full name of the skill I am using.	5 (7.4%)
		✓ The skill can read the full name configured for my Alexa account	61 (89.7%)
		✗ None of these.	0 (0%)
		I don't know.	1 (1.5%)
Mobile Number	69	✗ The skill will send me a email.	13 (18.9%)
		✓ The skill will know my phone number configured for my Alexa account	47 (68%)
		✗ The skill may access my phone calling history.	8 (11.5%)
		✗ None of these.	1 (1.5%)
		I don't know.	0 (0%)
Location Service	68	✗ The skill will ask me whether I want to play a voice message.	6 (8.8%)
		✓ A skill can know my dynamic location information.	47 (68.1%)
		✗ Allow a skill to access my device address.	11 (16.1%)
		✗ None of these.	3 (4.4%)
		I don't know.	2 (2.9%)
Email	69	✗ The skill will know my home address.	14 (20%)
		✓ A skill can read the email configured for my Alexa account	48 (69.5%)
		✓ The skill might send me a email.	7 (10%)
		✗ None of these.	0 (0%)
		I don't know.	0 (0%)
Amazon Pay	68	✓ Amazon Pay will share my name, email and shipping address	22 (32.3%)
		✓ Allow a skill to use Amazon Pay to make my payments.	40 (58.8%)
		✗ Someone will call me.	0 (0%)
		✗ None of these.	3 (4.4%)
		I don't know.	1 (1.5%)
Alexa Notification†	68	✗The skill will call me.	1 (1.5%)
		✗ Alexa will play a message from the skill without asking me.	20 (29.4%)
		✓ The skill will ask me whether I want to play a voice message.	36 (53%)
		✗ None of these.	7 (10.3%)
		I don't know.	3 (4.4%)
Alexa Reminder†	69	✗ The skill will ask me whether I want to play a voice message.	32 (46.3%)
		✓ Alexa will play a message from the skill without asking me.	22 (31.8%)
		✗ The skill can read my browser history.	13 (18.9%)
		✗ None of these.	0 (0%)
		I don't know.	2 (2.9%)
Read List†	68	✗ The skill may play music automatically.	20 (29.4%)
		✗ Allow a skill to modify my shopping list.	16 (23.5%)
		✓ The skill can read my shopping list.	25 (36.7%)
		✗ None of these.	5 (7.3%)
		I don't know.	2 (2.9%)
Write List†	69	✗ The skill will notify me about my bills.	16 (23.2%)
		✗ Alexa will play a message from the skill without asking me.	15 (21.7%)
		✓ The skill may change my shopping list.	30 (43.4%)
		✗ None of these.	4 (5.8%)
		I don't know.	4 (5.8%)

†: Four permissions were considered to be Alexa-specific. Note: the options for all these questions were shuffled for each survey.