Who am I Talking to? A Large-Scale Measurement of Surface Attribution Across Real-World Security and Privacy Interfaces

Marian Harbach
Google
Munich, Germany
mharbach@google.com

Jessica Johnson
Google
Mountain View, California, USA
johnsonjj@google.com

Abstract

Modern user interfaces are complex composites, with elements originating from various sources, such as the operating system, apps, a web browser, or websites. We posit that security and privacy decisions can to some extent depend on users correctly identifying an element's source, a concept we term "surface attribution." Through two large-scale vignette-based surveys ($N = 4,400$ and $N = 3,057$), we present the first empirical measurement of this ability.

We find that users struggle, correctly attributing UI source only 55% of the time on desktop and 53% on mobile. Familiarity and strong brand cues are associated with improved accuracy, whereas UI positioning, a long-held security design concept especially for browsers, has minimal impact. Furthermore, simply adding a "Security & Privacy" brand cue to Android permission prompts failed to improve attribution. These findings demonstrate a fundamental gap in users' mental models, indicating that relying on them to distinguish trusted UI is a fragile security paradigm.

CCS Concepts

• **Information systems** → *Web applications*; • **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → *Empirical studies in interaction design.*

Keywords

user interfaces, usable security, usable privacy, understanding, control, mental models

ACM Reference Format:
Marian Harbach and Jessica Johnson. 2026. Who am I Talking to? A Large-Scale Measurement of Surface Attribution Across Real-World Security and Privacy Interfaces. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 26 pages. https://doi.org/10.1145/3772318.3791287

1 Introduction

The software systems and corresponding user interfaces (UIs) that users have to interact with are quite complex. When using apps on mobile devices, users will be exposed to interface elements owned by the operating system, the handset manufacturer, third-party libraries as well as the app developers themselves. When people use websites and web apps, the browser is yet another source of user interface elements to consider.

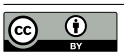

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791287

Ideally, users wouldn't have to care about who is showing which parts of the UI on their screen, so they are not unnecessarily burdened with technical details and fulfilling their needs is as effortless as possible. Yet, security and privacy mechanisms can rely on trusted UI, such as the line-of-death concept in browsers [9]. Such UIs ask users to make decisions, for example about permissions, or to enter sensitive information, like passwords. We posit that understanding the provenance of a supposedly trusted UI can be a key factor when getting phished or scammed, allowing a malicious piece of software to further embed itself into one's system, or simply knowing that there is a corresponding setting to go back on a previous decision.

A second, important aspect is users' general understanding of how the systems they use work, commonly referred to as mental models. We posit that understanding who is showing which parts of the UI is a key influence on having useful mental models of the systems' overall behaviors and thus feeling more confident in their use. It seems easier to derive what is happening in a complex UI if a question asked, option offered, or consequence observed can be associated with the correct actor. For example, when a user is asked whether or not they want to allow access to their camera while believing it is an app asking for this access, they might take this question as being asked essentially out of politeness or to manage expectations by the app. Instead, if they knew the question came from the operating system, they may be more likely to perceive this as the OS keeping the app in check, understand it as a privacy protection mechanism, and expect to configure this permission in the system settings rather than the app's. While the influence of mental models on human-computer interaction in general and security and privacy decision making in particular has been extensively researched for decades (e.g., [2, 17]), the influence of users' ability to understand which part of the stack of software they are interacting with in a given moment has not been explored before to the best of our knowledge.

In this paper, we present an initial exploration of a concept that has heretofore not been explored directly in the usable security space, to the best of our knowledge: the attribution of UI surfaces to the correct entity from the stack of software contributing parts of the overall user interface on the screen. In this first exploration, we intentionally focus on honest presentations of such UIs, given that any spoofing attempts by malicious actors will be easier or harder depending on how well users are able to correctly attribute the legitimate surfaces. We also do not attempt to link surface attribution to security and privacy behaviors. Instead, we focus on understanding the status quo. Based on two large-scale, vignette-based surveys (N=4,400 and 3,057), we provide insights on users' surface attribution, addressing three research questions: (1) To

what extent can users attribute browser- or OS-provided UIs to the correct entity? (2) Which factors contribute to correct attribution by users? and (3) To what extent does adding visual cues to Android permission prompts help to improve correct attribution?

We find that our sample of Chrome users can only correctly identify the source of a piece of UI in 53% of cases on mobile devices and 55% of cases on desktop or laptop computers. Key factors influencing correct attribution include surface familiarity and brand cues (such as logos and styling). Additionally, the presence of user data within the UI and whether the UI is triggered by a user action appear to have a more minor influence. UI positioning, however, was found to have minimal impact. Furthermore, simply adding "Security & Privacy" branding to Android permission dialogs didn't improve correct attribution either. Adding such branding did, however, help participants on non-Samsung devices to realize where they need to go in order to later change the corresponding setting.

Overall, our findings indicate that while visual cues are associated with a better understanding of who is providing a given piece of UI, just adding any branding will not be an instant remedy. Instead, we suspect that users have to get used to these cues through repeated exposure within and outside of their usage journeys. Hopefully, an increased understanding of UI surface attribution can lead to clearer user interfaces and thus help users to build better mental models and increase their control over the system. Finally, we caution that visual cues can also be a double-edged sword, as they are easy to fake, just like positive security indicators in general. To reduce risk in platform and system design, we believe that reducing the reliance on trusted UI for security and privacy purposes overall is the most meaningful path forward.

2 Background and Related Work

For users to make sound security and privacy decisions, they must often rely on cues from the UI. The provenance of these UI elements – whether they originate from the operating system, an app, the browser, or the content of a website – is critical, as they are not equally trustworthy. This work builds on a rich history of research in usable security and privacy, particularly in the areas of user mental models, trusted UI, and anti-phishing research. Yet, the question of users' ability to identify who is responsible for which parts of a composite user interface has not been investigated directly to the best of our knowledge.

2.1 Mental Models in Security and Privacy

A user's mental model is their internal understanding of how a system works, which they use to predict its behavior and guide their interactions [2, 11]. In the context of security, flawed mental models can lead to catastrophic errors. For instance, studies have consistently shown that many users possess inaccurate mental models of fundamental web security concepts, such as the meaning of the browser lock icon and the scope of HTTPS protections. Many users mistakenly believe the lock icon signifies a website's legitimacy or safety from all threats, rather than just a secure connection [16, 18].

This gap between the system's actual function and the user's understanding is a recurring theme in usable security and privacy research for almost two decades. Dhamija et al. [3] find that phishing works because users don't look at or are confused by which security UI is relevant. Jackson et al. [6] showed that browser-in-browser attacks are effective, likely due to confusion as to which is the "real" browser. In 2005, Ye et al. [21] already proposed to create "a trusted path from the browser to the human user", but this concept never came to fruition. The inability to attribute UI to its correct source is a fundamental breakdown in a user's mental model, preventing them from understanding who they are interacting with – a prerequisite for making an informed decision. Similarly, modern usable privacy research still highlights a need "for users to understand the inner workings of these interconnected processes to be able to inform themselves and, ultimately, make informed decisions" [19].

2.2 The Fragility of Trusted UI

Many security mechanisms depend on the concept of a Trusted User Interface (TUI) – a part of the screen that is assumed to be under the exclusive control of a trusted entity, such as the operating system (OS), browser, or a trusted website. Permission prompts, password entry dialogs, and payment forms are all examples of TUIs. They are the designated areas where users are expected to make critical decisions, from granting access to their camera to entering financial credentials.

However, the efficacy of this model rests on two fragile assumptions: 1) that the trusted entity can effectively prevent untrusted actors (like a malicious website) from spoofing the TUI, and 2) that the user can correctly identify the TUI and distinguish it from untrusted content. Decades of research on phishing and UI spoofing attacks show that both assumptions are frequently violated. Attackers have become adept at creating pixel-perfect replicas of legitimate login pages, browser chrome, and system dialogs [3, 7], effectively fooling users into disclosing sensitive information. Many solutions for phishing threats have been proposed [5], but most deployed measures focus on awareness and prevention: trying to educate users and flagging known phishing sites with systems like Google's SafeBrowsing[1]. This work investigates the second assumption of TUIs in a best case scenario: even when a UI is legitimate, can users correctly attribute it?

2.3 Aiding Attribution: Indicators and the Line of Death

Recognizing this challenge, the security community has proposed various solutions to help users distinguish trusted UI from web content. Early efforts focused on security indicators in browser chrome, such as the lock icon for secure connections. However, follow-up studies demonstrated that users often failed to notice or correctly interpret these indicators, and attackers could easily mimic them to create a false sense of security [16, 20]. The introduction and eventual deprecation of Extended Validation (EV) certificates, which displayed the organization's name in the address bar, further underscored the difficulty of creating a universally understood and effective visual indicator [6, 13].

The concept of the "line of death" – the visual boundary separating the trusted browser chrome from the untrusted web content – was hypothesized as a key demarcation that could help users make this distinction [9]. The theory suggests that UI appearing above

[1] https://safebrowsing.google.com/ – last accessed: 2025-07-01.

this line is trustworthy, while content below is not. Floating dialogs and prompts should always overlap the line of death, signaling to the user that this UI is impossible to draw for a website attempting to spoof it. However, this concept has been criticized for being too subtle for most users and for breaking down with modern web capabilities and UI designs [14]. The Chromium security considerations for browser UI [15] even state that "maintaining an understandable and consistent distinction between browser UI and web content is more of an art than a science", listing several principles to avoid the most devastating attacks. Our study empirically tests the effectiveness of spatial positioning, among others, as an attribution cue.

While much research has focused on the failure of specific indicators or the broad problem of phishing, there has been less focus on systematically quantifying users' fundamental ability to understand responsibility across a wide variety of composed UI elements in modern computing environments. This paper aims to fill that gap by providing a foundational, large-scale empirical measurement of this ability and identifying the factors that influence it.

2.4 Broader Context and Design Guidance

The challenge of surface attribution also connects to broader fields of research in website credibility, foundational UI design principles, and the guidelines provided to developers to create consistent user experiences.

Research on website credibility has long established that users make rapid judgments about a site's trustworthiness based on rather superficial design elements. Factors such as a professional and appealing visual design, the absence of typographic errors, and the presence of expected elements (like a privacy policy link) act as heuristics for perceived credibility [4, 8]. The findings of our study add a critical dimension to this research: if users cannot reliably determine the source of a trust-conferring UI element – for instance, whether a "safety badge" is provided by the browser or the website – then their credibility assessments are built on a fragile foundation. An attacker can easily spoof these visual heuristics, making a malicious site appear credible to a user who struggles with surface attribution.

Furthermore, the user confusion documented in our work can be thought of as a challenge with core principles of user interface design. Nielsen's heuristics, for example, emphasize "visibility of system status" and "consistency and standards" [10]. However, when a user cannot tell whether a dialog is from the operating system or the website because both parts of the UI leverage consistent and standardized components, the system's status is fundamentally unclear. Furthermore, while users might expect that elements from the same source will look and behave consistently, actors with malicious intent exploit this expectation, misleading users to believe that a given UI is from a legitimate, trustworthy source, while it is really part of the attack. In that sense, the challenge of surface attribution is exacerbated by these design principles establishing orthogonal goals.

Finally, platform owners like Apple, Google, and Microsoft invest heavily in creating and promoting design systems (e.g., Material Design[2], Fluent Design[3]). A primary goal of these guidelines is to create a consistent look and feel for components, both at the system as well as the application level. However, our findings suggest that these efforts are possibly too successful, as a harmonized look and feel across all UIs on a platform might actually hamper an understanding of who is responsible for which parts. Moreover, some common developer practices, such as creating "pre-prompts" that mimic the appearance of operating system (OS) dialogs, further muddy the waters for users. Our work provides empirical evidence of a tension between the intended simplicity and clarity of platform design systems and the real-world confusion experienced by users.

3 Methodology

In an ideal world, all users would understand who is responsible for each part of the user interface they are seeing at all times, so they can leverage this understanding to make better security and privacy decisions. We designed our studies to understand to what extent this ideal of correct surface attribution is realized today and what may be influencing users' understanding. Specifically, we set out to answer the following three research questions:

- RQ1. To what extent can users attribute browser- or OS-provided UIs to the correct entity?
- RQ2. Which factors contribute to correct attribution?
- RQ3. To what extent does adding visual cues to Android system prompts help to improve correct attribution?

As mentioned in the introduction, users are faced with UIs composed from many different sources in many of their typical user journeys. On mobile devices, such UIs would typically be displayed one after the other, with an app taking over the entire screen after the user selects it from their home screen. When engaging in a sensitive action during app use, users might be faced with a large, modal message from their operating system, for example asking for permission. In other cases, they might see an app surface next to an OS surface, like the status bar at the top of the screen, showing for example that location access is ongoing, or the keyboard, which might offer choices like selecting which password to fill into a login form. The set of passwords to fill might be provided by yet another app, like a password manager.

In contrast, on desktop devices with their larger screens, users are more likely to see UIs of differing provenance at the same time. Several windows of different applications might be next to or on top of each other. Operating system dialogs often appear as non-modal windows that may not be interacted with for a while, and other kinds of applications may show small, always-on-top messages from the OS status bar. Given this difference in UI interaction, we designed our studies to address mobile and desktop experiences individually.

The key dependent variable for this work is "correct attribution", meaning that participants can correctly identify who is responsible for a particular part of the UI from a list of possible entities (see questionnaires in the appendix). Attribution is a multi-faceted concept, ranging from who is responsible for the pixels being shown on the screen (responsibility), to where can the user change related

[2]https://material.io – last accessed 2025-07-01.
[3]https://fluent2.microsoft.design/ – last accessed 2025-07-01.

settings (management) and who caused the surface to show up (initiation). We posit that an understanding across these facets of attribution is central to establishing useful mental models, which then support users' informed decision making on security and privacy questions. For example, if users don't realize that permission prompts are provided by the operating system, they may also not realize that this choice provides them a real opportunity to protect their privacy when faced with an app asking to access their geolocation. If they instead believed that it is the app asking, they might not feel that they have a real choice as they may lack trust that the app would abide by their decision or they might not know which settings surface to go to if they need to revisit their decision later. Similarly, if users cannot distinguish between a prompt for their credit card details from their browser mediating a real payment flow and a website trying to phish them, they might either abandon real shopping journeys because of a lack of trust or submit sensitive information because of unjustified trust. However, to the best of our knowledge, this link between surface attribution and security and privacy behavior has not been investigated before. With this work, we aim to establish the status quo of surface attribution as a foundation for additional investigation.

For RQs 1 and 2, we designed and fielded a vignette-based survey that directly compares how well mobile and desktop users can attribute responsibility of surfaces shown by the Chrome browser on Windows and Android operating systems respectively. We focus on the Chrome browser on these platforms, as it presents the most common, yet more complex situation for users in everyday usage. Chrome is the most popular browser[4], while Android and Windows are the most popular platforms[5] at the time of designing the study. Furthermore, we chose to investigate a browser, as the browser is yet another source of UI surfaces on top of the OS and the app/website. In our second study, we address RQ1 with a focus on differences between mobile operating systems as well as RQ3 for a specific and important use case, granting permission for access to sensitive capabilities. This second study was centered around a mobile app asking for access to the user's microphone. We chose this context for study 2 to investigate if we have a chance to influence user perceptions in a simpler situation, with only two primary sources of UI surfaces. With study 2, we also explore differences in attribution across its various facets. We will describe both studies' methodologies below.

3.1 Study 1: Browser Surface Attribution

This study was commissioned through Ipsos, providing access to their panel and managing the fielding of the survey. The authors were responsible for questionnaire design as well as analysis in collaboration with Ipsos. The study targeted adults over 18 years of age in the US, focusing on a single, well-studied country for this first exploration. The study was conducted brand-blind, to avoid attracting respondents with overly positive or negative attitudes towards our organization.

Ipsos recruited a general population sample of respondents that needed to use Chrome on a Windows desktop or laptop computer or on their Android phone at least on a weekly basis. This restriction was necessary as we prepared screenshots for each UI to test and wanted the screenshots to closely resemble the platform respondents are most familiar with. We aimed to collect at least 200 responses per surface to allow for potentially investigating sub-group effects within each surface. Ipsos did not disclose the compensation each participant received but stated that they use industry-standard rates. The survey had an incidence rate of 29% and ran in July 2024. Participants were assigned to either the mobile or desktop group, based on which devices they indicated to be using on a regular basis. If they indicated to be using Chrome on both Windows and Android, they were randomly assigned to one of the groups. We collected a total of 4,400 complete responses with a median survey duration of 15 minutes. Participant demographics can be found in Table 1.

Table 1: Participant demographics for study 1.

Property	Value	Mobile	Desktop
N		2,024	2,376
Age	18-24	9%	9%
	25-34	15%	18%
	35-44	24%	23%
	45-54	25%	22%
	55-65	27%	28%
Gender	Female	49%	55%
	Male	51%	45%
Ethnicity	White	66%	67%
	Black or African-American	12%	11%
	Hispanic	12%	11%
	Asian/Pacific Islander	3%	5%
	Native American/Alaskan Native	2%	2%
Education	Education through Grade 12	3%	2%
	High-school graduate	20%	14%
	College	66%	64%
	After Bachelor's degree	11%	20%
Employment	Employed	64%	70%
	Unemployed	35%	29%
Marital Status	Married/Living w/ partner	57%	60%
	Single	27%	26%
	Divorced	14%	11%
	Widowed	3%	2%

3.1.1 Ethical Considerations. Our work was not subject to IRB review. Instead, a cross-functional team of stakeholders as well as user experience (UX) researchers at our organization reviewed and approved the research plan. All of the UX researchers involved in the project received formal training on research ethics.

Furthermore, we did not receive any identifying data beyond answers to our questionnaire from Ipsos. Participation in the survey was offered to volunteers subscribed to Ipsos' panel. Participants were informed about the nature of the survey and that this survey was commissioned by a client of Ipsos, while being offered an option to end their participation. When asking about sensitive demographic properties such as ethnicity, we specifically informed participants about how their data would be handled before the question and provided a "Prefer not to answer" option. Additionally, participants were able to end their participation at any point during the survey.

[4]https://gs.statcounter.com/browser-market-share, last accessed 2025-06-25.
[5]https://gs.statcounter.com/os-market-share, last accessed 2025-06-26.

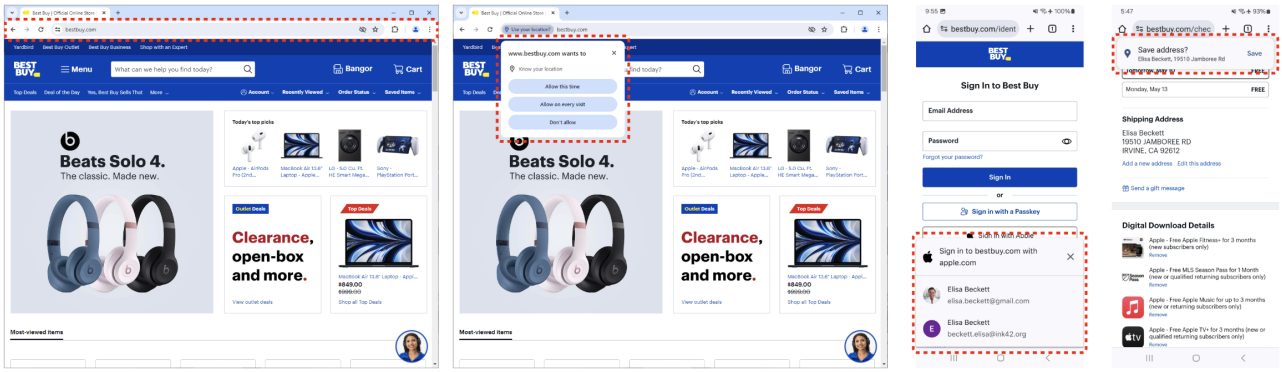

Figure 1: Four examples of the surface stimuli included in the study. The relevant surface is highlighted with a dashed red line. From left to right, the surfaces are desktop-control-topchrome, desktop-permission-plain, mobile-fedcm-other, mobile-autofill-save.

3.1.2 Survey Design. We collected a total of 36 desktop and 25 mobile surfaces for inclusion in this survey from Chrome version 126 by consulting with the Chrome team. They provided a list of surfaces that covers Chrome functionality that has a relationship to privacy and security. As a kind of control condition, we also included a number of neutral, commonly used Chrome surfaces, such as the address bar or the three-dot menu, as well as some surfaces that are not provided by Chrome at all, such as a website's search box or an OS-level permission prompt. The non-control surfaces varied along five key dimensions:

- **brand cues**: whether or not the surface contains any brand logos (Chrome or Google) or images that conform to Google's design language and color scheme;
- **containing user data**: whether or not the surface contains data from the user's profile, e.g. a username, credit card number, or address;
- **triggering**: whether or not the surface appears automatically or after a user action. The provided scenarios described when the surface appears;
- **position**: where the surface appeared, for example overlapping the address bar, near the address bar, or centered on the website content. This dimension was used to investigate the efficacy of the "line of death" concept; and
- **familiarity**: in contrast to the other dimensions, familiarity was assessed through a question in the survey (Q17 in Appendix B).

Some surfaces were slightly modified versions of the existing surfaces. For permission prompts, we included a version where the prompt was moved a few pixels down, so that it wasn't overlapping the address bar anymore. For FedCM[6] prompts, we included a version that did not use Google as the identity provider to introduce another brand. For each included surface, we provided a description of a situation in which a user would typically encounter it. We aimed to provide consistent scenarios across surfaces to the extent possible. Most surfaces were presented in a scenario where respondents were told to imagine going to bestbuy.com, a well-known e-commerce site in the US. Where this shopping-related scenario didn't fit, we introduced other situations, such as installing a Progressive Web App (PWA) on https://squoosh.app or clicking on an ad for headphones. You can find the list of tested surfaces alongside a description and the scenario provided to participants in Appendix A, the screenshots we used in the supplemental material, and several examples of surfaces in Figure 1.

The questionnaire was structured as follows. We first asked for consent and asked screening questions. After a participant qualified and was assigned to their group (mobile vs. desktop), they saw instructions for the surface attribution questions. The description outlined that participants will see several screenshots with a certain area highlighted in each one. We reinforced that we are only interested in who is responsible for the highlighted part of the screenshot. Then, each participant was randomly assigned to one control and three non-control surfaces. The control surface also appeared at a random position among the four total surfaces each participant saw. For each surface, participants were first asked to carefully review the scenario and screenshot, before proceeding to the next screens that asked them to indicate who they think is responsible for this surface. Finally, we asked additional questions about their browser use, tech affinity, as well as additional demographic questions. You can find the full questionnaire in Appendix B, which was refined in several iterations of pre-testing with colleagues.

3.1.3 Data Analysis. Given that each participant contributed data to four conditions each, we use Generalized Linear Mixed Models (GLMM) to ascertain whether differences in attribution are statistically significant, entering the respondent identifier as a random effect. We ran separate models for the desktop and mobile groups, given that the set of conditions is not the same and the screenshots are substantially different as well. For RQ1, we dummy-coded the attribution variable to test whether the proportion of a given response option is different between conditions, using the Chrome address bar as the reference level ('control-topchrome'). For RQ2, we ran one model per platform, using the five dimensions the surfaces varied across as independent variables and correct attribution

[6]https://developer.mozilla.org/en-US/docs/Web/API/FedCM_API – last accessed 2025-07-01

as the dependent variable. We use the *lme*4 package in R version 4.5.1 with the BOBYQA optimizer [12], as suggested by the lme4 authors [1]. When reporting percentages of correct attribution that differ from the overall percentage, we add the significance level and odds ratio from the model output.

3.2 Study 2: Mobile OS Permissions Attribution

In study 2, we focused on permission prompt UIs as a security and privacy measure commonly encountered by users to see if adding a visual indicator can improve surface attribution (RQ3). In contrast to study 1, we chose to conduct study 2 on mobile OS permission prompts instead of browsers to also gather additional data points for RQ1. Additionally, as noted above, browser permission prompts are an additional source of UI on top of the OS's UIs. To avoid any additional influences from this complexity, and thus give the additional visual indicator the best chance of working, we focused on OS-level permissions in this study. As our study is about user interfaces, we included the three most common mobile OS UIs for permissions: stock Android, Samsung's variation of Android, as well as Apple.

This study engaged participants through the Ugam Solutions panel, with surveys programmed by one of the authors in Qualtrics. Ugam was responsible for project management, including fielding the survey and initial data scrubbing. The study targeted adults across three distinct geographies: the US, UK, and India, aiming for approximately 1,000 complete responses from each country. To increase sample heterogeneity and representativeness and to more robustly test the presence of any effect beyond a single, culturally homogeneous group, we also implemented quotas based on participants' views on online privacy and their level of technology knowledge. Participants were recruited based on their reported phone type to ensure representation across the three common mobile operating system UIs. While our ideal approach would have been to recruit an equal number of participants for each group, we had to work within the practical constraints of the vendor's panel and our recruitment needs. In contrast to study 1, which employed a more general population pool to establish a baseline for broad-scope web interaction and attribution, study 2 was specifically designed to compare more nuanced interactions across distinct mobile operating system ecosystems. We also chose to exclude individuals working in the tech industry from study 2. By excluding tech workers, we ensure any effects we find pertain to a broader consumer audience. The specific quotas we set were a result of the vendor's recruiting abilities and results from prior work internal to our organization. While the exact compensation amount was not provided by Ugam, participant payment was reported to be around $5, consistent with standard rates for panel participation. The survey was fielded in March 2024 and the incidence rate for the survey was at least 20% in all three markets, with a total of 3,057 complete responses. Participant demographic information is summarized in Table 2. The median survey duration was 15 minutes.

3.2.1 Ethical Considerations. Similar to study 1, this work was not subject to IRB review. The research plan underwent review and approval by a cross-functional team of stakeholders and UX researchers within our organization, all of whom had received formal training on research ethics. No identifying participant data beyond survey responses were received from Ugam. Participation was voluntary for individuals subscribed to Ugam's panel. Prior to participation, individuals were informed about the nature of the survey, its commission by a client, and were provided the option to withdraw at any point.

Table 2: Participant demographics for study 2.

Property	Value	Overall
N		3,057
Gender	Female	48.7%
	Male	51%
	Non-binary	0.4%
	Other	0.2%
	Prefer not to say	0.1%
Age	18-23	19%
	24-30	18%
	31-40	20%
	41-50	20%
	51-60	18%
	60+	5%
Country	US	34%
	UK	33%
	IN	33%
Phone Type	Android (non-Samsung)	44%
	Samsung	43%
	Apple	14%
Privacy Views	Unconcerned	8%
	Fundamentalist	26%
	Pragmatist	66%
Tech Savviness	Less Tech Savvy	28%
	Tech Savvy	72%
Phone Model	Samsung	42%
	Motorola	11%
	OnePlus	4%
	Pixel	6%
	LG	2%
	Other	21%
	Apple (iOS Total)	14%

3.2.2 Survey Design. The study employed a between-subjects design to investigate attribution of mobile OS permission requests across variations of permission UIs and the impact of added security and privacy (S&P) branding cues. Participants were divided into three main groups based on their primary phone usage: an Apple group, a Samsung group, and an Android (non-Samsung) group (comprising users from various other Android OEMs like Pixel, OnePlus, Motorola, and LG). This grouping was established to examine potential differences in attribution given the distinct UI characteristics of Samsung devices and the broad adoption of design principles from Google's 'near-stock' Android experience by other Android OEMs.

The screener was structured as follows. After a welcome and instructions, participants first proceeded through a screener. This section ensured participants were not working in the tech industry, were at least 18 years old, and owned an eligible mobile phone model. Finally, participants were asked additional demographic questions about their gender, age, and mobile phone usage, as well as their views on online privacy and their knowledge of technology.

Next, participants were presented with a common scenario: "Imagine the following scenario: You download a new shopping app.

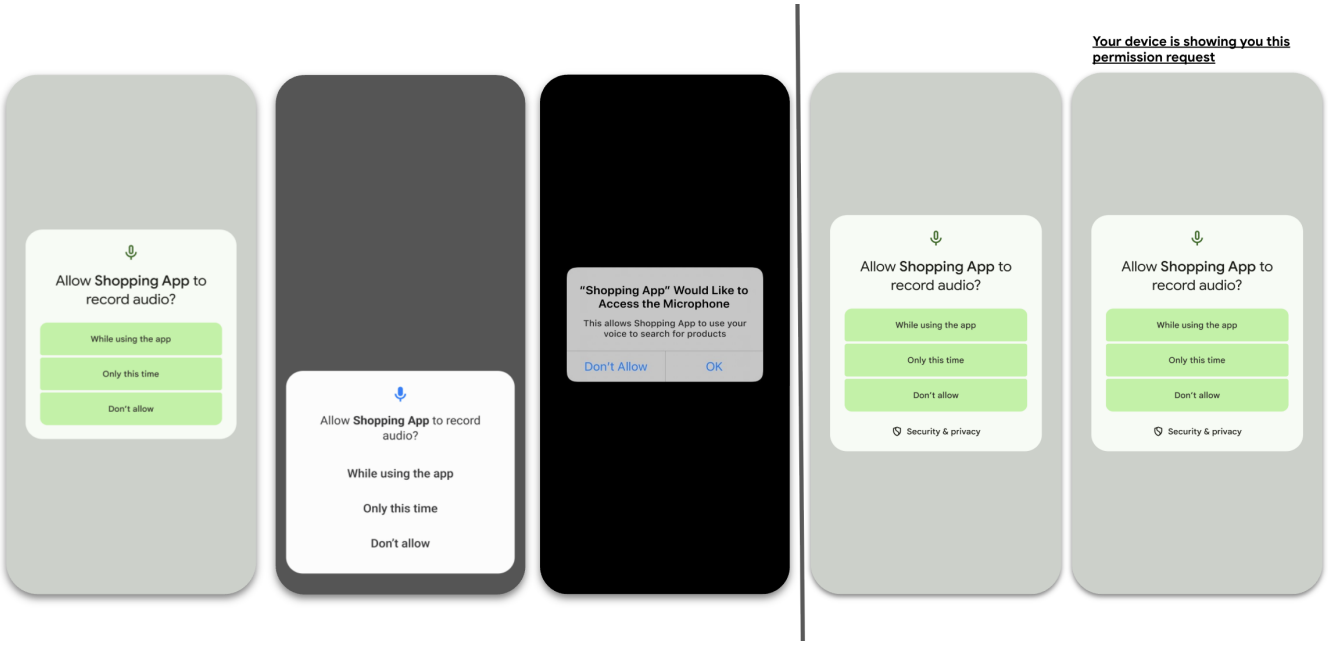

Figure 2: From left to right, the stimuli for Android (non-Samsung), Samsung, and Apple in the control group. The images on the right of the line show the two treatment conditions for the Android (non-Samsung) group, which we refer to as "S&P branding" and "S&P branding + explanatory text".

So, you open the app for the first time and get prompted with the following:". Following this, participants were shown a screenshot-like image of a permission dialog requesting microphone access. We chose microphone access as our permission type because it provided a simpler, more controlled variable for our study. Unlike location permissions, which often include additional safety labels and toggles for precise versus approximate tracking, the microphone permission request is straightforward. This allowed us to ensure that the only change in our experiment was the manipulated branding, avoiding any distractions that could have influenced participant responses. To maintain consistency across all conditions, "Shopping App" was chosen as a neutral application name. The specific appearance of this dialog varied by the participant's assigned phone group and experimental condition. You can find the full questionnaire in Appendix D.

To address RQ3, we created a new branding. We chose a plain black and white shield icon to symbolize protection, alongside a text reading "Security & privacy." We chose this branding as it could establish a clear, unifying thread that could plausibly help users correctly attribute protection moments to the device platform rather than the app across many different UIs. In contrast to the more obvious brand cues from Google and Chrome in study 1, the branding we tested for Android needs to work for many different Android OEMs, which is why we made it more neutral. We refer to this as S&P branding throughout the paper.

Control Group (without S&P branding): This group saw the standard permission prompt. The Android (non-Samsung) control group saw the standard Android UI with typical microphone icon at the top. The Samsung control group saw the identical text and options as Android (non-Samsung) Control, but presented with Samsung's distinct UI. Finally, the Apple control group saw the standard Apple UI (cf. Figure 2).

Treatment Group 1 (with S&P branding): For Android (non-Samsung) and Samsung groups, this condition added the S&P branding, placed at the bottom of the permission dialog (see second panel from the right in Figure 2). The Apple group did not have a corresponding treatment condition and served only as a baseline comparison in the control group.

Treatment Group 2 (with S&P branding + explanatory text): This "stress test" condition for Android (non-Samsung) and Samsung built upon Treatment Group 1 by including additional explicit text: "Your device is showing you this permission request," positioned directly above the mock permission dialog (see rightmost panel in Figure 2). This text was intended to maximize the potential impact of the S&P branding by explicitly clarifying its meaning in the context of this study, attempting to override potentially incomplete mental models. We showed this text outside of the mock to maximize the chances participants would actually notice and read it.

Each participant was exposed to only one of the permission dialog variants, consistent with the between-subjects design. Following exposure to the permission dialog and scenario, participants answered a series of questions. To comprehensively assess participants' surface attribution across facets, we employed several different operationalizations in this study to triangulate effects. These included:

Action-based attribution of control. Participants were asked, "If you selected an option and then wanted to change your selection at

another point in time, what would you be likely to do first?" (with options for "app settings" vs. "phone settings"). This question aimed to probe their understanding of where control over the permission resided.

Perceived attribution of responsibility and intent. Participants provided responses to questions about "which entity wants access to your microphone" and "which entity put this permission request pop up on your screen." These questions sought to understand their perception of the primary entity behind the prompt.

Perceived likelihood ratings for various entities. Participants rated the likelihood regarding the involvement of various entities (e.g., the app, the phone operating system, the phone manufacturer) in presenting the permission request. This provided a more quantitative measure of their attribution across plausible entities.

Prior to full data collection, we ran a pre-test on our survey questions with a small convenience sample (n=7) to ensure clarity and optimal wording.

3.2.3 Data Analysis. Given the between-subjects design, various statistical tests were employed to assess differences across experimental conditions and phone types. Dependent variables comprised participant responses to the attribution questions (e.g., "Open my phone's settings and look for the settings for this specific app" as the correct response for managing permissions; or attributing responsibility to "My device/OS" versus "the app" for surfacing the request). Independent variables included the assigned experimental condition (Control, Treatment 1, Treatment 2) and the participant's phone type group (Android (non-Samsung), Samsung, Apple).

Statistical analyses ran in SPSS included Pearson Chi-square tests and Comparisons of Column proportions to examine categorical attribution responses across groups. Paired-sample t-tests were conducted to compare the means of the likelihood ratings within conditions. Descriptive statistics and binomial tests were used to summarize and test specific proportions where appropriate. The Control group served as the primary baseline for comparisons across treatment groups.

4 Results

We will describe the results from the two studies as they pertain to our research questions in this section, before discussing the implications in Section 6.

4.1 RQ1: To what extent can users attribute browser- or OS-provided UIs to the correct entity?

Study 1 revealed that even though most surfaces were actually provided by Chrome, respondents frequently chose other attributions (cf. Figure 3). Overall only 55% and 53% of respondents were able to correctly identify who was responsible for showing the surfaces we tested on desktop and mobile respectively.

The most commonly correctly identified surfaces included password, payment, and authentication-related surfaces. Permission- and installation-related surfaces fared worst (cf. Table 3). For these surfaces, respondents much more frequently selected that the visited website is responsible (Desktop: $p < .001$ and OR between 3.2 and 9.7; Mobile: $p < .001$ and OR between 2.1 and 4.4). For the installation surfaces, respondents also more frequently selected "I don't know" (Desktop: $p < .001$ and OR between 7.6 and 10.9; Mobile: $p < .001$ and OR between 4.0 and 8.7). You can find a breakdown of attribution by surface in Tables 6 and 7 in Appendix C.

Table 3: The top 5 correctly and incorrectly attributed surfaces split between desktop and mobile.

Top 5			
Desktop surface	% correct	Mobile surface	%correct
password-save	82%	password-save	88%
help-customize	76%	payment-save	85%
fedcm-button	75%	password-suggest	83%
payment-save	74%	password-modal	80%
payment-modal	74%	passkey-save	77%

Bottom 5			
Desktop surface	% correct	Mobile surface	%correct
permission-context	32%	permission-context	28%
permission-bookmarks	29%	permission-plain	27%
help-me-write	26%	password-legacy	26%
permission-nocross	25%	install	20%
install	17%	permission-os	16%

Study 2 found similar results across the various operationalizations of attribution. We'll look at the results from the control groups in this section.

For the action-based attribution of control, significant differences were observed across phone types. Apple users demonstrated a significantly higher rate of correct attribution than Android (non-Samsung) users ($\chi^2(2) = 16.3, p < .001$). Specifically, 77% of Apple users correctly identified "Open my phone's settings" as the first step, compared to 64% of Android users and 70% of Samsung users (cf. Figure 4). These findings suggest that user behavior, platform-level UI design, and the permission prompt displays themselves may contribute to these differences in understanding where to manage permissions.

For perceived attribution of responsibility and intent, participants generally showed higher levels of correct attribution, though a notable proportion still made incorrect choices. For Android (non-Samsung) users, 55% correctly attributed the pop-up to the OS/device, while 45% incorrectly attributed it to "An app," "I'm not sure," or "Other." Samsung users had 62% correct attribution and Apple users also showed 62% correct attribution. Apple and Samsung users had significantly higher rates of correct attribution than Android (non-Samsung) users, but the difference is only 7 percentage points ($\chi^2(2) = 6.997, p = .030$) (cf. Figure 5).

Furthermore, on ratings for perceived likelihood of involvement in presentation of the permission prompt for various entities, Android (non-Samsung) users (average score of 3.5 on a scale where higher is more likely) showed a higher inclination to attribute it to the app, compared to their rating for the OS/device (average 2.7 for "Your device"). Samsung users showed a similar trend (3.4 for "The app" vs. 2.8 for "Your device"), as did Apple users (3.5 for "The app" vs. 3.3 for "Your device") (all $p < .005$). Notably, the Apple users' average likelihood rating for "Your device" (3.3) was significantly higher than that of both Android (non-Samsung) and Samsung users for their respective device categories ($p < .001$).

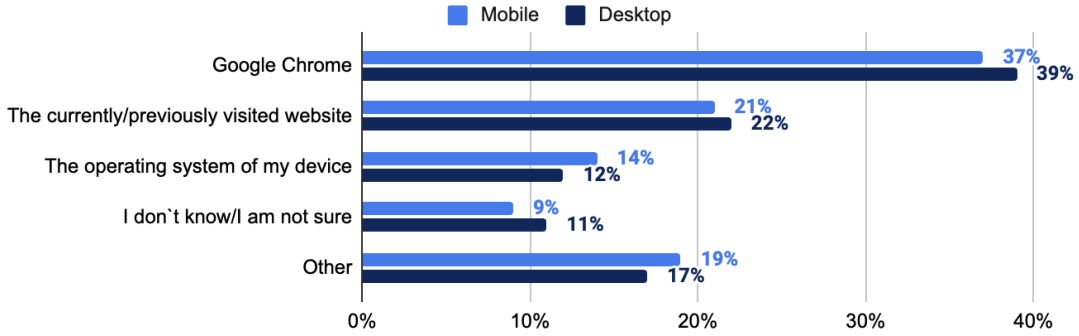

Figure 3: Overall responses to the attribution question (*"To the best of your knowledge, who is responsible for what you see in the highlighted area?"*).

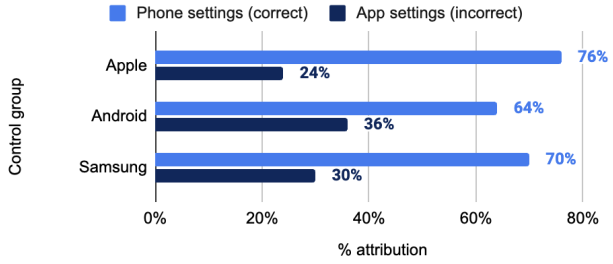

Figure 4: Overall responses to the attribution question (*"If you selected an option and then wanted to change your selection at another point in time, what would you be likely to do first?"*) in the control condition of study 2.

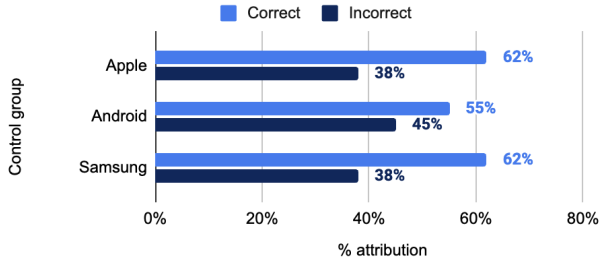

Figure 5: Overall responses to the attribution question (*"Which entity put this permission request pop up on your screen?"*) in the control condition of study 2.

Conversely, regarding "Which entity wants to access your microphone?" (where "An app" was the correct answer), participants showed significantly higher rates of correct attribution vs. incorrect attribution ($p < .001$ across all groups). 79% of Apple and Android (non-Samsung) users correctly identified the app as wanting access and 75% of Samsung users (see Figure 6). Differences between groups were not statistically significant ($\chi^2(2) = 3.49, p = .175$). This suggests that when it comes to the entity requesting the permission, Apple and Android (non-Samsung) users have similar understanding. Similarly, when asked to rate the likelihood that "the app" wants to access their microphone, average scores were consistently high across Apple (4.4), Android (non-Samsung) (4.3), and Samsung (4.2) users. No statistically significant differences were observed between the operating systems for any of the perceived entities wanting to access the microphone.

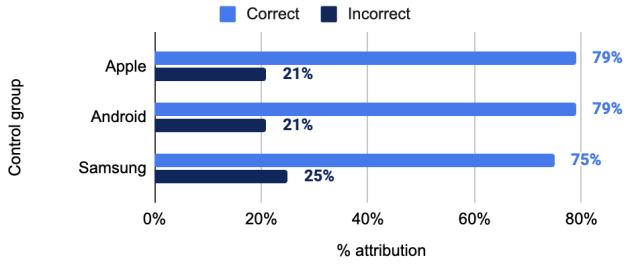

Figure 6: Overall responses to the attribution question (*'Which entity wants to access your microphone?'*) in the control condition of study 2.

Overall, while users are generally good at understanding that an app wants to access the microphone, there's a challenge in correctly attributing which entity (the OS/device versus the app) is displaying the permission request, and where to go to manage these settings. These findings are similar to what we found in study 1, but also suggest that attribution might be slightly easier for users of Apple. Additionally, levels of correct attribution were higher for mobile OS permission prompts in comparison to permission prompts in Chrome, which suggests that the additional layer of indirection when using the web could increase the challenge of attribution for users.

4.2 RQ2: Which factors contribute to correct attribution?

To explore which aspects influence correct attribution, we looked into each of the five dimensions across which the surfaces varied in study 1. Table 8 in Appendix C lists the results of the two GLMMs described in Section 3.1.3. We'll detail the influence of each of the dimensions in the following subsections.

Surface Familiarity. Familiarity varied substantially across surfaces (cf. Tables 6 and 7 in Appendix C). Between 27% and 79% of respondents indicated to be *'very'* or *'extremely familiar'* with the presented surface. The GLMMs we ran indicated that respondents rating a surface as 'very' or 'extremely familiar' were about twice as likely to attribute the surface correctly ($p < .001$, for both, mobile and desktop groups). Across all surfaces, correct attribution was higher by 14pp on desktop and 15pp on mobile when indicating that one is familiar with the given surface (cf. Figure 7).

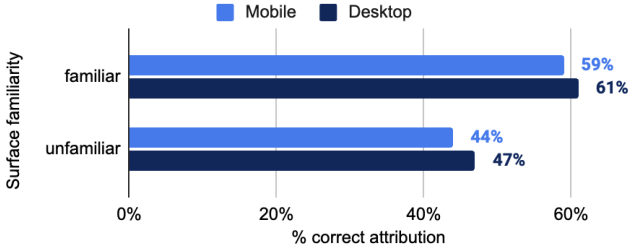

Figure 7: Correct attribution split by those respondents who rated themselves as 'very' or 'extremely familiar'.

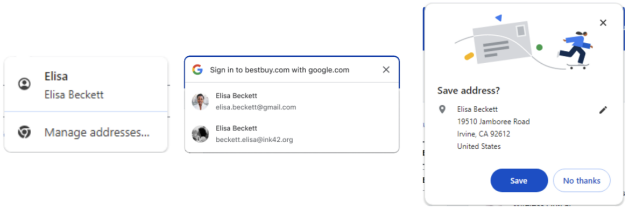

Figure 8: Three surfaces with the three different types of brand cues. The left-most surface has a Chrome logo, the surface in the center has a Google logo, and the right-most surface has an illustration using Google's style and colors.

Brand Cues. We assigned each surface to the dominant brand cue present on the surface, either the Chrome logo, the Google logo, or an illustration that uses Google's style and colors (see Figure 8 for examples). Seven mobile and 17 desktop surfaces contained brand cues. If any of the three types of brand cues were present, correct attribution improved between 13pp and 34pp (see Figure 9 – Desktop: OR Chrome logo 1.4, Google logo 2.3, illustration 1.7; Mobile: OR Chrome logo 1.3, Google logo 6.5, illustration 6.3; all $p < .001$). Interestingly, the increase is somewhat less pronounced on desktop in general and for desktop surfaces with an illustration in particular. It seems plausible that the larger surface area of the screen might make such brand cues less effective to grab the users' attention in comparison to the smaller screens of mobile devices.

Containing User Data. For desktop and mobile surfaces respectively, 8 and 6 surfaces included in our study mention user data (such as email address or first and last name), 10 and 6 include the website name, and another 5 and 4 had both. Surfaces containing user data were not associated with a significantly different level of correct attribution (see Figure 10).

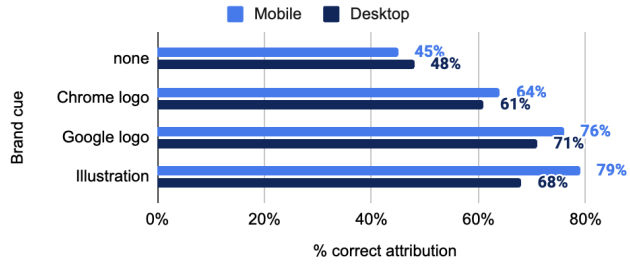

Figure 9: Correct attribution split by the type of brand cue the surface contained.

On the other hand, surfaces that only mention the website name (primarily permissions-related surfaces as well as PWA installation and the address bar) were associated with decreased correct attribution (-20pp and -21pp; Desktop: OR .34; Mobile: OR .64; all $p < .001$).

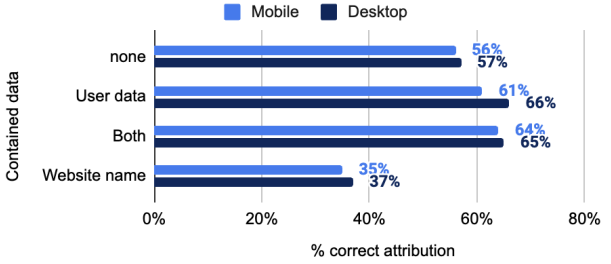

Figure 10: Correct attribution split by the type of data the surface contained.

Triggering. Our sample of surfaces contained 26 desktop and 18 mobile surfaces that typically occur immediately after the user did something. In comparison to those that do not follow a user interaction, surfaces triggered after an interaction exhibited an increase in correct attributions on mobile (cf. Figure 11; +12pp, OR 1.7, $p < .001$). This effect is also smaller than those we found for familiarity and brand cues. On desktop, the regression model suggests a weak negative effect on attribution (OR .84; $p < .05$).

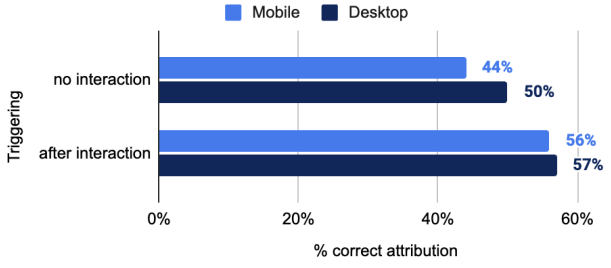

Figure 11: Correct attribution split by when the surface is triggered.

Positioning. In terms of positioning, our survey contained 14 desktop and 18 mobile surfaces that did not overlap the top bar. On desktop, we split further between 4 surfaces that are just barely below the address bar by a few pixels and another 10 surfaces that are further away or centered on the website content (such as autofill suggestions or certain variations of the permission prompt). For the surfaces that do overlap, we distinguish those that overlap only by a few pixels ('low' overlap, 17 desktop, 4 mobile) and those that have a more substantial overlap ('high', 7 desktop, 3 mobile). On mobile, high overlap includes the red malware warning, as it colors the entire address bar red. On desktop, we included the sidebar surface in the high overlap group, as it is visually connected to the address bar.

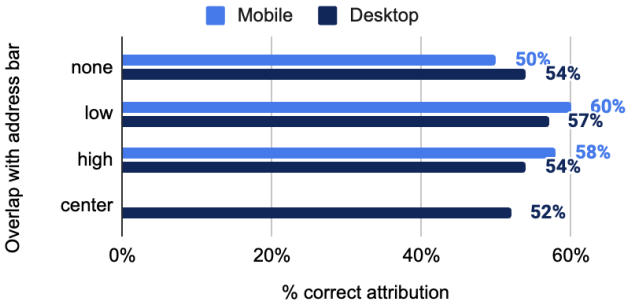

Figure 12: Correct attribution split by the extent the surface overlaps the browser's address bar.

As shown in Figure 12, there is a small positive effect on correct attribution based on surface positioning. On desktop, the GLMM shows a positive effect for low overlap surfaces (+3pp, OR 1.7, $p < .001$) and high overlap surfaces (+0pp, OR 1.5, $p < .001$). On mobile, the GLMM finds that both, low (+10pp, OR 2.4, $p < .001$) and high overlap (+8pp, OR 3.3, $p < .001$) have a significant positive effect. On desktop, the limited or non-existent differences in the aggregate percentage of correct attribution suggests that overlap is confounded with other factors. On mobile, the screen size as well as the type of UIs that use the overlap might make this distinction somewhat more understandable to users. We further explore the influence of positioning for permission prompts, a subset of the tested surfaces with otherwise similar properties, in the next subsection.

Deep Dive: Positioning of Permission Prompts. Study 1 contained seven variants of permission prompts for Chrome on desktop platforms[7]. Between these variants, only familiarity (between 58% and 76% rating themselves as very or extremely familiar), positioning (3 low overlap, 1 high overlap, 1 no overlap, and 2 centered), and triggering (4 not following user interaction, 3 following user interaction, cf. Table 5) varied. All seven variants had low levels of correct attribution (between 25% and 37%, cf. Table 6). Running the same GLMM on just these conditions (cf. Table 4), we find that neither triggering nor positioning retains a significant effect. Only familiarity is still associated with a positive effect (OR 1.85, $p < .001$). We further ran pairwise χ^2 tests between the permission prompt variants that only differed in positioning (legacy, plain, bookmarks, nocross, pepc-default, and pepc-noscrim on desktop) and did not find any significant effects. This supports that, at least for this particular surface, positioning does not significantly impact correct attribution.

Deep Dive: Pairwise Comparisons. Beyond permission prompts, a few other surfaces of the same kind also only differed in just one attribute. For one, payment-modal and payment-save on desktop as well as autofill-fill and autofill-save on mobile also only differ in positioning. We did not find any statistically significant differences in correct attribution across these pairwise comparisons. Second, fedcm-login-google and fedcm-login-other on both platforms only differed on the brand cue, one showing a Google logo and the other showing an Apple logo. For these conditions, we find a significant difference in correct attribution (Desktop: 69% vs. 37% correct, $\chi^2(1) = 41.6$, $p < .001$, Mobile: 67% vs. 33% correct, $\chi^2(1) = 61.9$, $p < .001$), congruent with the effects in the overall model.

Table 4: Impact of familiarity, positioning, and triggering on correct attribution for permission prompt conditions on desktop Chrome. The 'sig' columns indicate if a given level is statistically significantly different (*< .05, **< .01, *< .001) from the respective reference level. Reference levels: Familiarity – Not at all, slightly, or somewhat familiar; Active trigger – false; Overlap – none.**

Attribute	Value	odds ratio	sig.
(intercept)		0.17	**
Familiarity	very or extremely familiar	1.85	***
Active trigger	true	1.02	
Overlap	centered	1.90	
	low overlap	1.26	
	high overlap	1.47	

4.3 RQ3: To what extent does adding visual cues to Android system prompts help to improve correct attribution?

With this research question, we investigated the efficacy of adding Security & Privacy (S&P) branding cues to Android permission prompts in improving user attribution in study 2. We explored whether brand cues alone or in combination with an explanatory text could enhance users' understanding of the entity responsible for both surfacing and requesting permissions.

Our findings reveal a nuanced picture, suggesting that the brand cue we tested by itself and in a single exposure is insufficient in addressing attribution challenges, while adding explanatory text may help to at least make users action-ready, particularly for Android (non-Samsung) users.

Attribution of Permission Settings (Action-Based Understanding). For action-based attribution of control, the S&P branding treatments showed varied impact between Android OEMs. For Samsung users, neither the S&P branding alone nor the S&P branding with

[7]This does not include the "permission-quiet" condition, as this UI only consists of a chip in the address bar and not a full permission prompt.

text ("Your device is showing you this permission request") yielded a significant difference in attribution responses compared to the control group ($\chi^2(2) = 0.5, p = 0.8$). In the control group, 70% of Samsung users correctly identified "Phone settings" as the place to manage permissions. Following exposure to the S&P branding alone or with text, correct attribution remained similar at 72% (cf. Figure 13). These non-significant differences suggest that for Samsung users, these specific branding treatments did not notably alter their understanding of where to manage permissions.

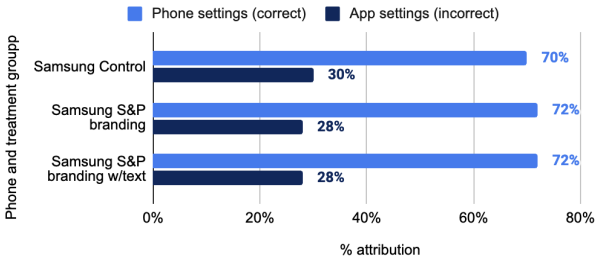

Figure 13: Overall responses to the attribution question (*"If you selected an option and then wanted to change your selection at another point in time, what would you be likely to do first?"*) in study 2.

In contrast, for Android (non-Samsung) users there was a significant overall difference in correct action-based attribution among the three groups ($\chi^2(2) = 13, p = .002$, cf. Figure 14). While the S&P branding alone did not make a significant difference in correct attribution in comparison to the control group, the S&P branding and explanatory text did demonstrably improve action-based attribution. In the control group, 64% of Android (non-Samsung) users correctly chose "Phone settings." With the S&P branding alone, correct attribution remained roughly the same with 66% correctly attributing. However, the S&P branding and text treatment significantly boosted correct attribution to 75%. This suggests that for these Android (non-Samsung) users, the explicit explanation was crucial in clarifying where to manage permissions. This improvement also brought attribution accuracy for managing settings on par with the Apple control condition (77% correct). This finding highlights the potential benefit of an additional, textual explanation.

Attribution of Entity Surfacing the Prompt. We examined which entity users believed was surfacing the permission request pop-up (correctly, the OS/device). This dimension revealed a crucial distinction between perceived understanding and action-based understanding, as users might choose a correct and helpful action, but their underlying beliefs reveal a different story.

For Android (non-Samsung) users, there was no statistically significant difference in correct attribution among the control group (55% correct), the S&P branding alone group (53% correct), and the S&P branding and text group (57% correct) ($\chi^2(2) = 2.087, p = .352$, cf. Figure 15). This suggests the S&P branding treatments did not substantially improve understanding. Within each of the three treatment groups (control, S&P branding, and S&P branding and explanatory text), we also compared the average likelihood rating for "the app" to "your device.". "The app" consistently received

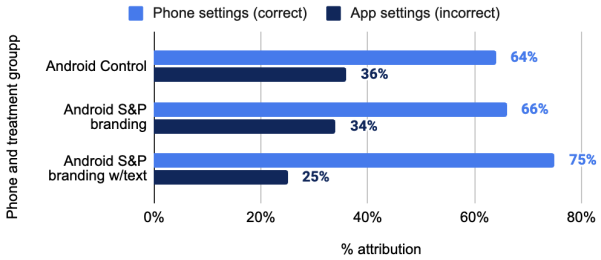

Figure 14: Overall responses to the attribution question (*"If you selected an option and then wanted to change your selection at another point in time, what would you be likely to do first?"*) for Android (non-Samsung) users in study 2.

significantly higher average responses than "your device" (control: 3.5 for "The app" vs. 2.7 for "Your device" $p < .001$; S&P branding: 3.4 vs. 2.6, $p < .001$; S&P branding and text: 3.4 vs. 2.9, $p < .001$). These findings indicate a persistent, fundamental misunderstanding of the permission system at a belief level, even while the action-based understanding shows some improvement.

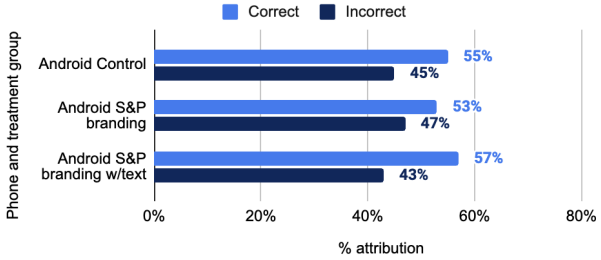

Figure 15: Overall responses to the attribution question (*"Which entity put this permission request pop up on your screen?"*) by Android (non-Samsung) users in study 2.

Similarly, for Samsung users, the S&P branding (61% correct) and S&P branding and text (64% correct) also showed no significant differences from the control (62% correct) in attributing which entity surfaced the prompt ($\chi^2(2) = 1.195, p = .55$, cf. Figure 16). Users continued to conflate the app's desire for permission with the OS/device's role in displaying the request. The ratings for "the app" were consistently higher than "your device" (e.g., control: 3.4 for "The app" vs. 2.8 for "Your device", $p < .001$; S&P branding: 3.5 vs. 2.9, $p < .001$; S&P branding and text: 3.5 vs. 2.9, $p < .001$), reinforcing this confusion.

While study 1 suggested that branding can contribute to improved attribution, the findings from this study show that simply adding a brand cue or explanatory text was insufficient to significantly improve participants' understanding of which entity surfaces the permission request on Android. Potential reasons could include that the brand cue we chose was not well-known or that it takes more than a single exposure for such an intervention to have an effect. This highlights the need for more research on the impact of interventions to clarify attribution.

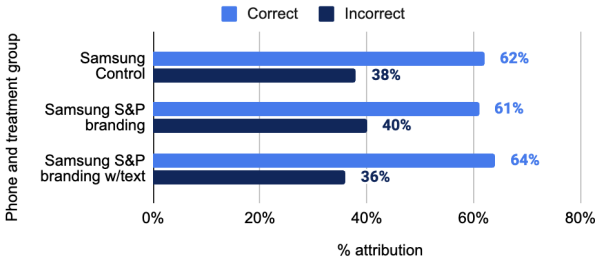

Figure 16: Overall responses to the attribution question (*"Which entity put this permission request pop up on your screen?"*) by Samsung users in study 2.

Attribution of Entity Requesting Permission Access. Finally, we examined which entity users believed was *requesting* microphone access (correctly, the app). Here, users already exhibited high levels of correct attribution in the control group. The S&P branding treatments had only limited impact across user groups.

For Samsung users, there was a significant difference among groups when responding to which entity requested microphone access ($\chi^2(2) = 6.217, p = .045$, cf. Figure 17). The S&P branding (79% correct) had a significantly higher correct attribution rate than both the control (75% correct) and the S&P branding and text (75% correct) groups. Yet, this effect is very small.

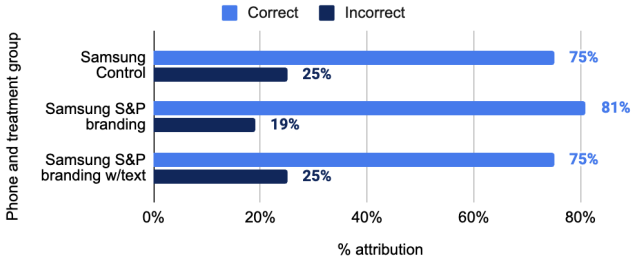

Figure 17: Overall responses to the attribution question (*'Which entity wants to access your microphone?'*) for Samsung users in study 2.

For Android (non-Samsung) users, the S&P branding and S&P branding with text description did not significantly impact their attribution for the entity requesting the permission. Android (non-Samsung) users showed consistent correct attribution rates (S&P branding: 79% correct; S&P branding and text: 80% correct), similar to the control condition, with no significant differences across treatment groups ($\chi^2(2) = 0.971, p = .059$, cf. Figure 18).

Likelihood ratings also reflected this strong understanding; "The app" consistently received much higher average scores across all Android (non-Samsung) and Samsung groups (e.g., for Android S&P branding: 4.2 for "The app" vs. 2.3 for "Your device", $p < .001$; for Samsung S&P branding: 4.3 for "The app" vs. 2.3 for "Your device", $p < .001$).

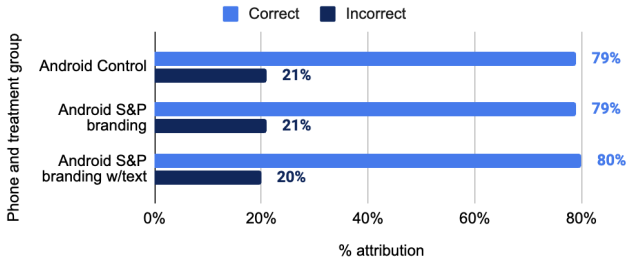

Figure 18: Overall responses to the attribution question (*'Which entity wants to access your microphone?'*) for Android (non-Samsung) users in study 2.

5 Threats to Validity

There are several threats to the validity of the findings in our two studies.

Ecological validity. Both of our studies used a synthetic survey setup. While we provided contextual information in the form of vignettes, looking at screenshots is still different from actually using a product for a purpose of one's own choosing. It stands to reason that additional signals about surface attribution might be available from a sequence of interactions or richer context. On the other hand, participants of online surveys might be paying more attention to the nuances visible in the presented surfaces, while users might only briefly glance at the same surface in a real-world situation. This could lead to responses that are not representative of how users would make decisions when engaged in an entirely separate primary use case. An experience sampling-based approach could help to overcome this challenge, but is of course more difficult to field.

Non-experimental study design. We chose the surfaces to include in study 1 based on advice from the Chrome team and exploration of available surfaces. While we identified several key dimensions across which the included surfaces varied, we did not ensure that all combinations of these dimensions are equally represented in our study, nor did we control for additional aspects that might affect respondents' judgments. Similarly, study 2 focused only on permission prompts, which are a common security and privacy UI users encounter. Yet, they are also only one particular example and other surfaces may yield different behavior changes when branding would be added. Thus, our results are more of an observational and exploratory nature and further work is needed to establish the observed patterns more rigorously with a more controlled study design on a broader set of situations. Finally, we intentionally focused on existing, honest UIs as opposed to an attacker's spoofing attempts to first establish how well users can reason about legitimate interactions. How susceptible users are to misdirection based on various spoofing variations is an important piece of future work.

Asking about attribution. Who is responsible for a given part of the UI on a screen is not a question that comes up naturally for most people. It is therefore challenging to ask such a question in a research study. We chose our question wording carefully, iterating on it in pre-testing, and even included a detailed description

and example of what we mean in study 1. Study 2 further operationalized attribution in several ways, triangulating respondents' understanding to some extent. Nevertheless, it is possible that our questions were still misunderstood and thus contributed to the levels of incorrect attribution we observed.

Additional threats. We exposed participants of study 1 to four surfaces in a somewhat repetitive task design. This could have led to fatigue, which we counteracted by randomization. Furthermore, we chose to field this exploratory study only in the US, as this population is relatively well understood due to the many usable security and privacy studies that have been conducted on similar samples. Future work should expand the scope to include other countries and cultures as well to establish patterns more robustly. Additionally, for study 2, our use of quotas based on prior internal work and the vendor's recruiting abilities could introduce selection bias. While we aimed for a representative sample, the specific quotas might not fully represent the general population of users in the target geographies. Finally, we chose to focus our investigation on the influence of properties of the UI surfaces on correct attribution. It seems very plausible that properties of users, such as digital literacy, would also mediate the understanding of complex UIs, which should be subject to a dedicated investigation.

6 Discussion

The findings from our two large-scale studies reveal a significant gap in the ability to correctly attribute the source of user interface (UI) elements, a concept we term surface attribution. With a correct attribution rate of only 53% on mobile and 55% on desktop, it seems very likely that users can often struggle to identify whether a UI element originates from the operating system, the browser, or a website. This potential for fundamental misunderstandings has profound implications for users' roles in security and privacy. Some safety mechanisms rely on users' ability to recognize and trust UI from a specific source. For example, UIs that ask for sensitive user information, such as passwords or biometric data, rely on users only entering this sensitive information on an appropriate, trusted UI surface. If users can't distinguish trusted from untrusted or spoofed versions of this UI, successful phishing attacks may become more likely.

Another set of privacy- and security-relevant UIs are those that ask users to make decisions, such as the permissions prompts investigated in our studies. We believe that misunderstandings about who is responsible for showing such UIs may lead to inaccurate mental models, which could then lead to unsafe behavior and confusion. As mentioned in the introduction, a user that does not understand that it is the OS or the browser that asks whether or not a website or app should have a given permission may be less careful about their decision due to a lack of perceived protection or may be less likely to find the associated settings. The same can be true for surfaces such as auto-fill and payments, which also manage privacy-relevant, sensitive information. We hope to substantiate this relationship between correct surface attribution and behavior in future work.

On a positive note, we find that a larger fraction of respondents in our second study correctly identified system settings as the place for changing their permission decisions despite lower levels of understanding that the OS is responsible for showing the permission UI. This finding hints at some disconnect between understanding what is happening and still being able to find the right place to configure an associated setting. Across the three variations of attribution we investigated in study 2, both action-based attribution and understanding who will get access to the capability showed higher levels of correct answers in comparison to attributing who is responsible for the content of the surface. This difference could be explained by the ability to learn about who gets access and where to make changes, based on normal device usage: users experience apps being able to use the microphone after a access was granted, and similarly, users may go into settings somewhat frequently and can thus see where permission settings are available. In contrast, understanding who is responsible for showing the decision UI is not explained anywhere, there are only subtle design clues users could pay attention to, and reasoning about surface attribution is not a common need for users to begin with. The differences we find suggest that future studies should look at attribution as a multifaceted concept, including action-based understanding as well as a more direct understanding of who is responsible for a given piece of UI to discern an understanding of consequences and possible remedies from a deeper understanding of the underlying system.

Our research also indicates that while some factors can improve surface attribution, there doesn't seem to be a simple solution. Notably, familiarity and the presence of brand cues like logos and specific styling was associated with substantial increases in correct attribution. At the same time, study 2 showed that adding a simple "Security & Privacy" branding to a permission prompt did only have a very limited effect. While action-based attribution was improved for Android (non-Samsung) users when paired with additional explanatory text, perceived understanding of who surfaced the prompt did not change. In our study, this explanatory text was presented to participants outside of the mobile user interface itself. This distinction is important: while our work demonstrates the cognitive benefit of such an explanation, it does not prescribe a specific UI implementation, nor does it account for the significant design challenges, such as UI clutter and potential alert fatigue, that would arise from integrating more text into permission dialogs. Therefore, this finding should not be interpreted as a simple recommendation to add more text, but rather as an illustration of the comprehension gap in mental models that needs to be addressed. How to best bridge that gap within the strict constraints of a mobile UI remains a complex and open question and requires further research. Yet, our findings still suggest that consistent exposure and clear branding can help users build more accurate mental models, but maybe only over time and when the branding has sufficient familiarity. This underscores that visual cues are not a panacea and can be a double-edged sword, as they are easily spoofed by malicious actors. Additional aspects worth exploring further include the presence of a user's own data within a UI element, which also slightly improved attribution, likely because it signals that a trusted entity with access to that data is responsible.

Conversely, the positioning of a UI element, a factor historically considered important in security design (e.g., the "line of death" concept), had a minimal impact on correct attribution in our study,

especially on desktop platforms with their larger screens. In a direct comparison on permission prompts, positioning did not significantly impact correct attribution. This finding challenges long-held assumptions in the security community and suggests that spatial cues alone are not enough to help users distinguish trusted UI from potentially malicious content. This aligns with modern critiques that argue such visual boundaries are too subtle for the average user to notice and interpret correctly.

6.1 Implications for Usable Privacy and Security Research

Our findings present several key takeaways and directions for future work for researchers and practitioners in the usable privacy and security space.

Rethinking Reliance on Trusted UI. The low overall proportion of correct surface attribution suggests that relying on users to distinguish trusted UI is a fragile security model. The community should continue to explore and advocate for security mechanisms that are less dependent on user perception and more on technical safeguards that work in the background.

The Power of Consistent Branding and User Education. While not a complete solution, the positive impact of familiarity and brand cues suggest their value in helping users understand. Yet, additional work should more rigorously establish the efficacy of such UI patterns with controlled experiments. Our exploration in study 2 only showed a limited effect.

Yet, if future work can demonstrate that the effects of branding and familiarity we found hold, platform and application developers could strive for consistent and recognizable UI for security-critical interactions. However, our findings also suggest that users need to be educated over time to recognize and correctly interpret these cues. Future research could also further investigate the effectiveness of different educational interventions in improving surface attribution. Finally, the reliance on visual cues also highlights that users remain susceptible to phishing and other spoofing attacks.

Beyond Visual Cues. The limited success of adding branding highlights the need for a more holistic approach. Instead of focusing solely on visual indicators, future work should explore multi-modal and interactive approaches to conveying the provenance of UI elements. This could include cues like haptic feedback for trusted interactions or system-level "lenses" that allow users to inspect the source of on-screen elements.

Investigating the Behavioral Impact of Misattribution. Our work establishes that users often misattribute UI elements. The critical next step is to understand the behavioral consequences of these errors. Future studies should investigate whether incorrect surface attribution directly leads to differences in permission decision-making, a higher susceptibility to phishing, malware installation, or other security and privacy harms.

7 Conclusion

In conclusion, the challenge of surface attribution is a fundamental issue with modern user interfaces that has been largely overlooked. Our research provides a first, foundational understanding of this problem, demonstrating that users' ability to identify the source of UI is far from guaranteed. Future work is needed to more firmly establish effects of surface attribution on mental models and user behavior. For the usable privacy and security community, our findings underline the need to not further burden the user with indicators, decision UIs, and warnings, and instead continue to move towards more robust, system-level security that reduces the complexity of today's user experiences.

Acknowledgments

Parts of the text in this paper have been refined using the Google Gemini LLM. The authors would like to thank Yu-Hsuan Lin, Hank H Huang, and Subin Shin for their valuable contributions to the branding design of the permission prompts. We also extend our gratitude to Laura Neiswander, Fabio Carnevale Maffe, and Rodrigo Farell for their support and contributions throughout the various stages of this project. Finally, we are thankful for the guidance provided by the anonymous CHI reviewers.

References

[1] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. doi:10.18637/jss.v067.i01

[2] John M. Carroll and Judith Reitman Olson. 1988. Mental Models in Human-Computer Interaction. In *Handbook of Human-Computer Interaction*, Martin Helander (Ed.). 45–65. https://www.sciencedirect.com/science/article/pii/B9780444705365500075

[3] Rachna Dhamija, J. D. Tygar, and Marti Hearst. 2006. Why phishing works. In *Proceedings of ACM CHI*. https://doi.org/10.1145/1124772.1124861

[4] B.J. Fogg. 2003. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann Publishers Inc.

[5] Anjuli Franz, Verena Zimmermann, Gregor Albrecht, Katrin Hartwig, Christian Reuter, Alexander Benlian, and Joachim Vogt. 2021. SoK: Still Plenty of Phish in the Sea — A Taxonomy of User-Oriented Phishing Interventions and Avenues for Future Research. In *Symposium on Usable Privacy and Security (SOUPS)*. USENIX Association. https://www.usenix.org/conference/soups2021/presentation/franz

[6] Collin Jackson, Daniel R. Simon, Desney S. Tan, and Adam Barth. 2007. An Evaluation of Extended Validation and Picture-in-Picture Phishing Attacks. In *Financial Cryptography and Data Security*, Sven Dietrich and Rachna Dhamija (Eds.).

[7] Markus Jakobsson and Steven Myers (Eds.). 2006. *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft*. John Wiley & Sons. https://onlinelibrary.wiley.com/doi/book/10.1002/0470086106

[8] Michal Kakol, Radoslaw Nielek, and Adam Wierzbicki. 2017. Understanding and predicting Web content credibility using the Content Credibility Corpus. *Information Processing & Management* 53, 5 (2017). https://www.sciencedirect.com/science/article/pii/S0306457316306471

[9] Eric Lawrence. 2017. The Line of Death. https://textslashplain.com/2017/01/14/the-line-of-death/. Last accessed: 2025-06-30.

[10] Jakob Nielsen. 1994. 10 Usability Heuristics for User Interface Design. https://www.nngroup.com/articles/ten-usability-heuristics/. Last accessed: 2025-06-30.

[11] Donald A. Norman. 1983. Some Observations on Mental Models. In *Mental Models*, Dedre Gentner and Albert L. Stevens (Eds.). Lawrence Erlbaum Associates.

[12] Michael JD Powell et al. 2009. The BOBYQA algorithm for bound constrained optimization without derivatives. *Cambridge NA Report NA2009/06, University of Cambridge, Cambridge* 26 (2009), 26–46.

[13] Jennifer Sobey, Robert Biddle, Paul C Van Oorschot, and Andrew S Patrick. 2008. Exploring user reactions to new browser cues for extended validation certificates. In *ESORICS: European Symposium on Research in Computer Security*. Springer.

[14] Emily Stark. 2022. The death of the line of death. https://emilymstark.com/2022/12/18/death-to-the-line-of-death.html. Last accessed: 2025-06-30.

[15] Chromium Team. [n. d.]. Security Considerations for Browser UI. https://chromium.googlesource.com/chromium/src/+/refs/heads/main/docs/security/security-considerations-for-browser-ui.md. Last accessed: 2025-07-01.

[16] Emanuel von Zezschwitz, Serena Chen, and Emily Stark. 2022. "It builds trust with the customers" - Exploring User Perceptions of the Padlock Icon in Browser UI. In *2022 IEEE Security and Privacy Workshops (SecWeb)*. https://ieeexplore.ieee.org/abstract/document/9833869

[17] Rick Wash. 2010. Folk models of home computer security. In *Symposium on Usable Privacy and Security (SOUPS)*. https://doi.org/10.1145/1837110.1837125

[18] Tara Whalen and Kori M. Inkpen. 2005. Gathering Evidence: Use of Visual Security Cues in Web Browsers. In *Proceedings of the Graphics Interface Conference (GI '05)*.

[19] Maximiliane Windl, Magdalena Schlegel, and Sven Mayer. 2024. Exploring Users' Mental Models and Privacy Concerns During Interconnected Interactions. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 259 (Sept. 2024). https://doi.org/10.1145/3676504

[20] Min Wu, Robert C. Miller, and Simson L. Garfinkel. 2006. Do Security Toolbars Actually Prevent Phishing Attacks?. In *ACM CHI*.

[21] Zishuang (Eileen) Ye, Sean Smith, and Denise Anthony. 2005. Trusted paths for browsers. *ACM Trans. Inf. Syst. Secur.* 8, 2 (May 2005), 153–186. https://doi.org/10.1145/1065545.1065546

A Study 1: Included Surfaces

Condition label	Surface Description	Scenario Description	Platform	correct	Overlap (D / M)	Cues (D / M)	Contains info	Active trigger
control-topchrome	Chrome's address bar including the navigation controls and the extensions and three-dot menu symbol.	You open a new browser window for bestbuy.com and you notice the area highlighted by the red dotted box	DM	C	high	none	website	✗
control-threedot	The menu that appears after clicking on the three-dot menu button in the address bar.	You're visiting bestbuy.com. You click on the three dots in the top-right corner of the Chrome window and the menu highlighted by the red dotted box appears.	DM	C	low / high	none	none	✓
control-contextmenu	The menu that appears after right-clicking on the content of a website.	You're visiting bestbuy.com. You right-click on the page and the area highlighted by the red dotted box appears.	D	C	center	none	none	✓
control-negative	The bestbuy.com search box with suggestions after typing "headphones".	You're visiting bestbuy.com. When you type "headphones" in the search box the area highlighted by the red dotted box appears.	DM	W	center / none	none	none	✓
filepicker	The OS-provided file picker dialog after clicking a button on squoosh.app.	You want to reduce the file size of an image on your computer. You open https://squoosh.app and click a button to upload a file. The area highlighted by the red dotted box appears.	DM	O	high / none	chrome / none	none	✓
install	The full PWA installation dialog shown by Chrome after clicking the install button.	You often visit the site https://squoosh.app. You click on the "Install" button on this page and then the popup highlighted by the red dotted box appears.	DM	C	low / none	none	website	✓
install-content	The developer-provided screenshots provided inside of the PWA installation dialog shown by Chrome after clicking the install button.	You often visit the site https://squoosh.app. You click on the "Install" button on this page and then the popup highlighted by the red dotted box appears.	DM	W	low / none	none	website	✓
permission-legacy	The standard Chrome permission prompt for geolocation access.	You're visiting bestbuy.com and immediately the area highlighted by the red dotted box appears.	DM	C	low / none	none	website	✗
permission-plain	The updated Chrome permission prompt with an "Allow only this time" option for geolocation access.	You're visiting bestbuy.com and immediately the area highlighted by the red dotted box appears.	DM	C	low / none	none	website	✗
permission-bookmarks	The same as the plain variant, but with the bookmarks bar visible for more overlap with the top part of Chrome.	You're visiting bestbuy.com and immediately the area highlighted by the red dotted box appears.	D	C	high	none	website	✗
permission-nocross	The same as the plain variant, but the permission prompt was moved several pixels down so that it no longer overlaps with the top bar.	You're visiting bestbuy.com and immediately the area highlighted by the red dotted box appears.	D	C	none	none	website	✗
permission-quiet	The quiet permission prompt UI Chrome shows for likely unwanted permission prompts as a chip on the left-hand side of the address bar.	You're visiting bestbuy.com and immediately the area highlighted by the red dotted box appears.	DM	C	high / low	none	none	✗
permission-context	The same as permission-plain, but the website in the background shows a store locator UI and the scenario mentions clicking a button.	You're visiting bestbuy.com and click on "Find store". On the "Find store" page, you click on "Use my current location" and then the area highlighted by the red dotted box appears.	DM	C	low / none	none	website	✓
permission-pepc-default	Similar to the permission-context condition, but the permission prompt appears centered over the website, while the website is being greyed out with a semi-transparent scrim.	You're visiting bestbuy.com and want to find the nearest store. You click on a button labeled "Use precise location" and then the area highlighted by the red dotted box appears and the background becomes gray.	D	C	center	none	website	✓
permission-pepc-noscrim	The same as permission-pepc-default, but without the scrim.	You're visiting bestbuy.com and want to find the nearest store. You click on a button labeled "Use precise location" and then the area highlighted by the red dotted box appears.	D	C	center	none	website	✓
permission-os	The permission prompt shown by the OS when geolocation is used for the first time.	You're visiting bestbuy.com and want to find the nearest store. You're asked if you want to share your location with bestbuy.com and you click Allow. Afterwards, the area highlighted by the red dotted box appears.	M	O	none	none	none	✓
red-interstitial	The warning page shown by Chrome when a page is likely malicious.	You click on an ad while searching for offers on a pair of headphones you want to buy. When the new site loads, the area highlighted by the red dotted box appears.	DM	C	none / high	none	none	✗
autofill-fill	The dropdown that appears when filling form fields.	You're shopping on bestbuy.com and want to finalize an order. When the site asks for your shipping information, you click into the "First Name" field and the area highlighted by the red dotted box appears.	DM	C	center / none	chrome / none	user	✓

Condition label	Surface Description	Scenario Description	Platform	correct	Overlap (D / M)	Cues (D / M)	Contains info	Active trigger
autofill-save	The prompt that appears after submitting a form with manually filled fields, asking to save the form information for later use.	You're shopping on bestbuy.com and just entered your shipping address to finalize your order. You click "Next", and the next page shows the area highlighted by the red dotted box.	DM	C	low	illu / none	user	✓
password-modal	A prompt that appears when Chrome can fill username and password.	You want to log in to bestbuy.com and just clicked into the 'username' field. The area highlighted by the red dotted box appears.	DM	C	low / none	none / illu	user	✓
password-fill	The dropdown that appears when filling sign-in forms.	You want to log in on bestbuy.com and just clicked into the 'username' field. The area highlighted by the red dotted box appears.	DM	C	center / none	chrome / none	both	✓
password-save	The prompt that appears after submitting a sign-in form, asking to save the login information for later use.	You just typed in your username and password on bestbuy.com to log in to your account. After clicking "Login", the page reloads and the area highlighted by the red dotted box appears.	DM	CP	low	illu / none	user	✓
password-suggest	The dropdown that appears when there is no account avialble in Chrome's password manager, suggesting to create a new strong password.	You want to create an account on bestbuy.com. After typing your name and email address, you click into the password field. The area highlighted by the red dotted box appears.	DM	CP	center / none	illu	user	✓
password-compromised	The warning prompt that appears when signing in with a password that has been detected as compromised in a data breach.	You want to log in to bestbuy.com. Chrome autofill just filled in the username and password fields. After clicking "Login", the page reloads and the area highlighted by the red dotted box appears.	DM	CP	low / none	illu	none	✓
passkey-save	The prompt that appears when creating a new passkey for a website.	You want to add a passkey to your account on bestbuy.com. After clicking the "Add passkey" button, the area highlighted by the red dotted box appears.	DM	CP	low / none	illu	both	✓
fedcm-login-google	The "Sign-in with ..." prompt that appears if a website uses the FedCM API[8]. Google is the identity provider in this case.	You want to log in to bestbuy.com. After navigating to the login page, the area highlighted by the red dotted box appears.	DM	C	none	google	both	✗
fedcm-login-other	The "Sign-in with ..." prompt that appears if a website uses the FedCM API. Apple is the identity provider in this case.	You want to log in to bestbuy.com. After navigating to the login page, the area highlighted by the red dotted box appears.	DM	C	none	none	both	✗
fedcm-button	The prompt that appears after clicking a "Sign-in with Google" button when a website is usign the FedCM API. The prompt offers to select which account to sign in with.	You want to log in to bestbuy.com. You click "Sign In With Google" and the area highlighted by the red dotted box appears.	D	C	low	google	both	✓
payment-save	The prompt that appears after entering credit card details and submitting the form, asking to save the payment details for later use.	You want to check out your shopping cart at bestbuy.com and just entered your credit card details. After you click next, the area highlighted by the red dotted box appears.	DM	CG	none	google	user	✓
payment-fill	The dropdown that appears when focusing a payment-related form field, offering to fill in credit card details.	You want to check out your shopping cart at bestbuy.com. After you click on the field for your credit card number, the area highlighted by the red dotted box appears.	DM	CG	center / none	chrome / none	user	✓
payment-modal	A prompt that appears to complete a checkout flow, for example when using Google Pay. The prompt contains an iframe from the payment provider.	You want to check out your shopping cart at bestbuy.com. After you click on a Google Pay button, the area highlighted by the red dotted box appears.	D	CG	low	google	user	✓
sidebar	The sidebar that opens after clicking "Search Image with Google" from the context menu.	You're visiting bestbuy.com. You right-click on an image of headphones and select "Search Image with Google". The area highlighted by the red dotted box appears.	D	C	high	google	none	✓
help-me-write	The prompt that appears when using Chrome's "Help me write" functionality.	You're visiting bestbuy.com. You want to write a review for a product and decide to use AI to help. You right-click on the text box and select "Help me write" from the menu. The area highlighted by the red dotted box appears.	D	C	center	none	none	✓
help-memory	A blue bubble that appears pointing at the three-dot menu, informing the user that Chrome has a Memory saver feature when the device is currently low on memory.	You're watching a video on youtube.com and the area highlighted by the red dotted box appears.	D	C	low	none	none	✗
help-customize	A blue bubble that appears centered and just below the address bar, informing the user that they can customize the appearance of Chrome.	You just restarted Chrome and opened a new tab. The area highlighted by the red dotted box appears.	D	C	low	none	none	✗
help-credit-card	A blue bubble that appears next to a credit card autofill dropwdown, notifying the user that this credit card is avaialble in autofill because it was saved to the Google Wallet associated with this account.	You're visiting bestbuy.com. You're checking out your shopping cart and clicked on the credit card number field. A suggestion for your credit card number as well as the area highlighted by the red dotted box appears.	D	C	center	chrome	none	✓

[8]https://privacysandbox.google.com/cookies/fedcm/why#user-interaction, last accessed: 2025-06-27.

Condition label	Surface Description	Scenario Description	Platform	correct	Overlap (D / M)	Cues (D / M)	Contains info	Active trigger
hats-invitation	The prompt that appears in the top right corner next to the three-dot menu inviting the user to participate in an in-product survey.	You're visiting bestbuy.com. You're looking for a BestBuy store nearby and were asked if you wanted to allow the site to use your location. You click 'Allow' and a few seconds later the area highlighted by the red dotted box appears.	DM	C	low	chrome	none	✓

Table 5: List of surfaces included in study 1 between mobile and desktop conditions. The platform column indicates whether this surface was present on desktop (D) and/or mobile (M). The correct column indicates which answer option was deemed correct during analysis (C=Chrome, G=Google Pay, O=Operating system, P=Password Manager, W=Website). The overlap column denotes to what extent the surface overlapped with the browser address bar (low=a few pixels, high=more than a few pixels, none=no overlap, center=centered on the content) and if this was different between the desktop and mobile versions. The cues column denotes whether there were any visual cues present on the surface (chrome=Chrome logo, google=Google logo, illu=illustration in Google/Chrome style, none=no visual cue). The 'contains info' column denotes if the surface contained at least one piece of information from the given source (website=the website's origin, user=user information, such as name, email address, or credit card info, none=no information present). Finally, the 'active trigger' column denotes whether or not the scenario description mentions a user action before the surface appears.

B Study 1: Questionnaire

Questions:

Q1. What is your date of birth?

- Year: 1910-2015
- Month: January - December

Q2. Are you ... ?

- Female
- Male
- Another gender
- Prefer not to answer

Q3. Please insert your zipcode: <write-in>

Q4. We at Ipsos protect our research and our client's confidential/proprietary information and intellectual property. By participating in this study, it is mandatory that you agree to: Not share any of the information included in this study such as images, videos, recordings advertisements and technical concepts, ideas, services, products or packaging; Not photograph, record, publish on the internet whatsoever, copy, or in any other way reproduce any of the information included in this study; Not use our client's confidential information for your own benefit, the benefit of a third party, or in any way which would negatively impact our client or their public image. Do you acknowledge you have read and agree to the statements above?

- Yes, I agree and I wish to participate in this study
- No, I do not agree

Q5. Which of the following devices do you currently own and regularly use in your household? (Select all that apply.)

- Computer/Laptop
- Smartphone
- Tablet (e.g., Apple iPad, Samsung Galaxy Tab, etc.)
- Smartwatch (i.e. used for calls, texts, notifications, fitness, and/or health)
- Fitness band / tracker (i.e. used primarily for tracking fitness and/or health)
- Smart home speaker (e.g., Amazon Echo, Google Home)
- Smart devices, other than speakers (e.g., smart lights, WiFi-connected cameras, doorbells, etc.)
- Smart TV or streaming device (has built in "apps")
- None of the above

Q6. What is the brand of your smartphone?
If you own more than one smartphone, please consider the smartphone that you use most often for personal reasons. (Select one.)

- iPhone/Apple
- Samsung
- Pixel/Google
- Motorola
- Nokia
- LG
- OnePlus
- Sony
- Other (Specify)

Q7. What is the operating system of your computer/laptop?
If you own more than one computer/laptop, please consider the computer/laptop you use most often for personal reasons. (Select one)

- macOS
- Windows
- Linux
- Chrome OS
- Other Operating System
- Unsure

Q8. Below is a list of web browser(s) you have heard of before. For each, please indicate which of the following best applies to you. (Select one for each.)
Please think about all your devices while answering the following.

- Google Chrome
- Apple Safari
- Microsoft Edge
- Mozilla Firefox

Answer Options

- Have not heard of or used
- Heard of, but have not used
- Used before but not currently
- Currently use this browser **only**
- Currently use this browser **in addition to others**

Q9. For each of the following devices you own, which web browser(s) do you currently use?
If you own multiple devices, please consider the device that you use most often for personal reasons. (Select all that apply for each device.)

- Computer/Laptop *(if using)*
- Smartphone *(if using)*

Answer Options: Pipe in Browsers currently used from Q8

Q10. How often do you use the following web browser(s) on your [INSERT DEVICE]? (Select one for each) *Statements: Pipe in browsers used on device from Q9 Answer Options:*

- Daily or more often
- A few times a week
- About once a week
- Monthly
- Less than monthly

Q11. Was your [INSERT DEVICE IN BOLD] provided to you by your employer, school, or another organization? (Select one)

- Yes
- No
- Not sure/ don't know

Q12. In this next section we will present you with various screenshots you may encounter while browsing the web on your [INSERT DEVICE]. For each example, we've highlighted a specific area with a red-dotted line to focus on.
We are interested in your assumption on who creates this part of the screenshot. We are not focused on why it is there or what caused it to show up. We want to know who you think created that specific piece of the screenshot and is responsible for the way it looks and what it says.
Desktop screenshot:

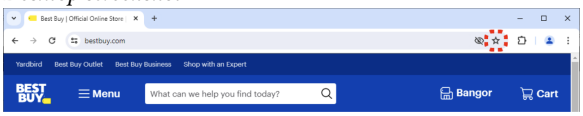

Mobile screenshot:

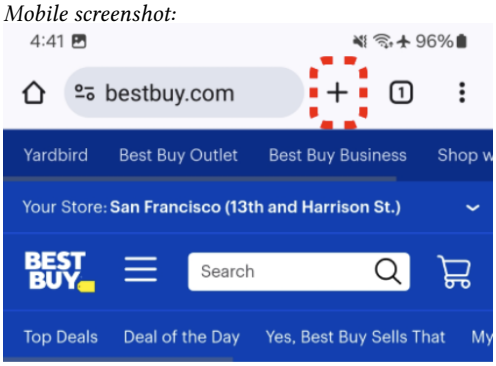

Loop Q13-Q18 for one control and 3 non-control surfaces

Q13. Please review the following scenario and screenshot carefully, paying close attention to the contents highlighted by the red-dotted outline.
Imagine you're using Chrome on your [INSERT DEVICE].
[INSERT SURFACE-SCENARIO IN BOLD ITALICS]
[DON'T LET RESPONDENT MOVE TO THE NEXT SCREEN UNTIL 10 SECONDS HAVE ELAPSED]

Q14. To the best of your knowledge, who is responsible for what you see in the highlighted area? (Select one)

- Google Chrome
- [INSERT SURFACE.WEBSITE WITH UPPERCASE FIRST LETTER] (i.e. the currently visited website) [IF SURFACE.WEBSITE != NA] [IF SURFACE.ID = RED INTERSTITIAL SHOW: 'The previously visited website']
- My internet service provider
- The operating system of my device (IF Computer/Laptop: 'Windows'; IF Mobile Android: 'Android'])
- The manufacturer of my device [SHOW IF Mobile]
- An extension installed on the device [SHOW IF Computer/Laptop]
- Google Search [SHOW IF SURFACE.ID = SIDEBAR]
- My admin at work, school, etc. [SHOW IF Q11='Yes']
- Google Password Manager [SHOW IF SURFACE.ID = PASSWORD]
- Google Pay [SHOW IF SURFACE.ID = PAYMENT]
- Other (Specify)
- I don't know/I am not sure

Q15. What leads you to believe that [INSERT Q13 RESPONSE BOLDED] is responsible for what you see in the highlighted area? (Please type your response below)

Q16. Which of the following aspects, if any, leads you to believe that [INSERT Q13 RESPONSE BOLDED] is responsible for what you see in the highlighted area? (Select all that apply)

- There is a brand logo
- Color scheme is recognizable
- Font style is recognizable
- Design elements (e.g., shapes, symbols) are recognizable
- The location of the highlighted area
- The brand was mentioned
- The type of feature/information being shown
- Feature works across all websites I visit
- Feature appears as a pop-up or dialog box
- The sequence of actions that leads to seeing the highlighted area
- I have seen and/or used before
- Other (Specify)
- None of the above

Q17. How familiar are you with seeing the highlighted area on your [INSERT DEVICE]? (Select one)

- Extremely familiar
- Very familiar
- Somewhat familiar
- Slightly familiar
- Not at all familiar

Q18. How frequently do you see the highlighted area on your [INSERT DEVICE]? (Select one)

- Daily
- Weekly
- Monthly
- Yearly or less frequently
- Never

Q19. Thank you for your answers so far! For the following questions, we will be asking you about your experiences with Google Chrome on your [INSERT DEVICE BOLDED].

Q20. What language are menu items, settings, options, etc. set to when using Chrome on your [INSERT DEVICE BOLDED]? (Select one)

- English
- Non-English
- None of the above

Q21. [IF LAPTOP] While using Chrome on your computer do you have the Chrome bookmarks bar showing / pinned? (Select one)

- Yes
- No
- Not sure/ don't know

Q22. Thank you for your answers! You are almost done, we just have a few more questions to learn more about you.

Q23. Which of the following best describes when you buy or try out new technology? (Select one)

- I'm among the first people to try
- I'm sooner than most people, but not among the first
- I use it once many people are using it
- I use it once most people are using it
- I am usually not interested in buying or trying out new technology

Q24. How would you rate your competency with technology and devices? (Select one)

- **Excellent**: I'm highly capable with technology and can easily troubleshoot problems myself
- **Very good**: I use technology in a variety of settings and applications; I require some help with troubleshooting
- **Acceptable**: I use technology with relative ease, but require a moderate amount of help with troubleshooting
- **Developing**: I'm not at ease or confident using technology; I require a lot of troubleshooting and/or assistance with technology from family and friends
- **Unfamiliar**: I'm not familiar with using technology and prefer not to use it

Q25. What is the highest degree or level of school you have completed? Select only one.

- Education through Grade 12 [Expandable Header]
- Grade 4 or less
- Grade 5 to 8
- Grade 9 to 11
- Grade 12 (no diploma)
- High School Graduate [Expandable Header]
- Regular High School Diploma
- GED or alternative credential
- College or Some College [Expandable Header]
- Some college credit, but less than 1 year
- 1 or more years of college credit, no degree
- Associate's degree (AA, AS, etc)
- Bachelor's degree (BA, BS, etc.)
- After Bachelor's Degree [Expandable Header]
- Master's degree (MA, MS, MBA, etc.)
- Professional degree (MD, DDS, JD, etc.)
- Doctorate degree (PhD, EdD, etc.)

Q26. What is your current employment status?

- Employed full-time
- Employed part-time
- Self employed
- Unemployed but looking for a job
- Unemployed and not looking for a job/Long-term sick or disabled
- Full-time parent, homemaker
- Retired
- Student/Pupil
- Military
- Prefer not to answer

Q27. Please indicate your annual household income before taxes.

- From < $5,000 to $250,000 or more
- Prefer not to answer

Q28. This is a topic of sensitive nature. Answering is voluntary, however, collecting such information enables us to provide a more refined research analysis. If you don't agree to provide us such information, a "Prefer not to answer" option is available for you to select, at your discretion. For any survey research purposes, your responses are combined with the answers from all other participants. We will provide our client only anonymous results.
Are you of Hispanic, Latino or Spanish origin? Select only one.

- Yes
- No
- Prefer not to answer

Q29. This is a topic of sensitive nature. Answering is voluntary, however, collecting such information enables us to provide a more refined research analysis. If you don't agree to provide us such information, a "Prefer not to answer" option is available for you to select, at your discretion. For any survey research purposes, your responses are combined with the answers from all other participants. We will provide our client only anonymous results.
What is your race? Select all that apply.

- White
- Black or African American
- Native American or Alaskan Native
- Asian
- Pacific Islander
- Other race
- Prefer not to answer

Q30. What is your marital status? Select only one.

- Single, never married
- Living with partner
- Married
- Widowed
- Divorced or separated

C Study 1: Attribution Details

C.1 Desktop group

Table 6: Attribution results from the desktop group. Columns show how many respondents chose this attribution option. The 'correct' column indicates how many of these choices we considered correct. The 'sig' columns indicate if this proportion is statistically significantly different (*< .05, **< .01, *< .001) from the reference level, control-topchrome. The value in brackets shows the odds ratio. The overall n value shows the unique respondent count. The 'familiar' column denotes which fraction of respondents rated this surface as 'very' or 'extremely familiar'.**

condition	n	correct	correct_sig	chrome	chrome_sig	website	website_sig	os	os_sig	idk	idk_sig	other	familiar
password-save	204	82%	*** (4.3)	27%	*** (0.2)	3%	*** (0.1)	5%	* (0.4)	8%		57%	75%
help-customize	204	76%	*** (3)	76%	*** (2.9)	11%	* (0.6)	4%	** (0.3)	5%		3%	49%
fedcm-button	204	75%	*** (2.6)	75%	*** (2.6)	7%	*** (0.3)	6%	* (0.4)	7%		5%	66%
payment-modal	204	74%	*** (2.5)	11%	*** (0.1)	7%	*** (0.3)	3%	** (0.3)	11%		67%	43%
payment-save	204	74%	*** (2.1)	37%	*** (0.3)	8%	** (0.4)	5%	* (0.5)	7%		43%	62%
password-fill	204	73%	*** (2.4)	26%	*** (0.2)	11%	* (0.6)	6%	* (0.5)	6%		51%	72%
help-memory	203	72%	*** (2.3)	72%	*** (2.3)	4%	*** (0.2)	10%		9%		5%	39%
control-negative	476	71%	*** (1.9)	13%	*** (0.1)	71%	*** (22.1)	3%	*** (0.2)	10%		3%	57%
hats-invitation	204	71%	*** (2.3)	71%	*** (2.2)	12%		3%	** (0.2)	9%		5%	45%
password-compromised	204	71%	** (2)	18%	*** (0.1)	7%	*** (0.3)	4%	** (0.3)	13%	* (2.6)	58%	42%
fedcm-login-google	204	69%	** (1.9)	69%	** (1.9)	11%		2%	*** (0.2)	12%		5%	66%
passkey-save	204	69%	* (1.7)	30%	*** (0.2)	11%		3%	** (0.2)	12%	** (3.6)	43%	41%
password-modal	204	65%	* (1.5)	28%	*** (0.2)	18%		5%	* (0.4)	6%		42%	61%
password-suggest	204	63%		22%	*** (0.1)	21%		5%	** (0.3)	9%		43%	69%
payment-fill	203	63%		28%	*** (0.2)	21%		4%	** (0.3)	9%		37%	57%
control-threedot	475	62%		62%		5%	*** (0.2)	22%	*** (2.6)	7%		4%	71%
sidebar	203	62%		30%	*** (0.2)	17%		5%	* (0.4)	7%		41%	41%
filepicker	475	58%		18%	*** (0.1)	8%	*** (0.4)	58%	*** (25.7)	11%		6%	69%
control-topchrome	**475**	**57%**	–	**57%**	–	**17%**	–	**11%**	–	**8%**	–	**6%**	**78%**
autofill-save	204	55%		55%		15%		8%		12%		10%	54%
autofill-fill	204	54%		54%		16%		11%		15%	* (2.8)	4%	60%
install-content	204	49%	** (0.5)	17%	*** (0.1)	49%	*** (6.9)	7%		20%	*** (10.9)	8%	30%
help-creditcard	204	49%	* (0.6)	49%	** (0.6)	14%		3%	*** (0.2)	12%		23%	40%
permission-quiet	204	49%	* (0.6)	49%	* (0.6)	25%	* (1.8)	5%	* (0.5)	11%	* (2.9)	10%	54%
red-interstitial	203	48%	** (0.6)	48%	** (0.6)	3%	*** (0.1)	19%	** (2.2)	12%	* (3.1)	19%	38%
control-contextmenu	475	38%	*** (0.3)	38%	*** (0.3)	7%	*** (0.3)	39%	*** (8.6)	11%		5%	68%
fedcm-login-other	204	37%	*** (0.3)	37%	*** (0.4)	17%		11%		18%	*** (6.9)	17%	42%
permission-pepc-default	204	37%	*** (0.3)	37%	*** (0.4)	43%	*** (4.3)	5%	* (0.4)	9%		6%	63%
permission-pepc-noscrim	204	36%	*** (0.3)	36%	*** (0.3)	33%	*** (3.3)	7%		14%	** (3.7)	10%	58%
permission-legacy	203	33%	*** (0.2)	33%	*** (0.3)	51%	*** (7.2)	5%	* (0.4)	5%		7%	76%
permission-plain	204	33%	*** (0.3)	33%	*** (0.3)	50%	*** (7.1)	4%	* (0.4)	9%		4%	69%
permission-context	204	32%	*** (0.3)	32%	*** (0.3)	42%	*** (4.7)	7%		11%	* (2.8)	8%	64%
permission-bookmarks	203	29%	*** (0.2)	29%	*** (0.2)	46%	*** (5.8)	6%		11%		8%	61%
help-me-write	204	26%	*** (0.2)	26%	*** (0.2)	34%	*** (3.2)	4%	** (0.3)	21%	*** (9.8)	15%	31%
permission-nocross	204	25%	*** (0.2)	25%	*** (0.2)	55%	*** (9.7)	6%	* (0.5)	10%		4%	65%
install	204	17%	*** (0.1)	17%	*** (0.1)	52%	*** (8.1)	4%	** (0.3)	15%	*** (7.6)	11%	28%
Overall	**2,376**	**55%**		**39%**		**22%**		**12%**		**11%**		**17%**	**57%**

C.2 Mobile group

Table 7: Attribution results from the mobile group. Columns show how many respondents chose this attribution option. The 'correct' column indicates how many of these choices we considered correct. The 'sig' columns indicate if this proportion is statistically significantly different (*< .05, **< .01, *< .001) from the reference level, control-topchrome. The value in brackets shows the odds ratio. The overall n value shows the unique respondent count. The 'familiar' column denotes which fraction of respondents rated this surface as 'very' or 'extremely familiar'.**

condition	n	correct	correct_sig	chrome	chrome_sig	website	website_sig	os	os_sig	idk	idk_sig	other	familiar
password-save	264	88%	*** (9.3)	31%	*** (0.4)	2%	*** (0)	3%	** (0.3)	5%	** (0.2)	59%	75%
payment-save	265	85%	*** (6.8)	21%	*** (0.2)	3%	*** (0.1)	5%	** (0.4)	5%	*** (0.2)	66%	62%
password-suggest	265	83%	*** (5.6)	29%	*** (0.3)	4%	*** (0.1)	4%	** (0.3)	7%	* (0.3)	56%	61%
password-modal	264	80%	*** (4.7)	28%	*** (0.3)	8%	*** (0.2)	6%	* (0.5)	4%	** (0.2)	55%	73%
passkey-save	264	77%	*** (3.8)	34%	*** (0.4)	8%	*** (0.2)	5%	* (0.4)	7%		46%	47%
password-compromised	264	76%	*** (3.5)	25%	*** (0.3)	5%	*** (0.1)	7%		8%		56%	41%
control-threedot	506	68%	*** (2.2)	68%	*** (2.2)	4%	*** (0.1)	16%	* (1.6)	6%	** (0.3)	7%	73%
fedcm-login-google	264	67%	*** (2.1)	67%	*** (2.1)	11%	*** (0.4)	10%		7%		4%	62%
hats-invitation	264	64%	*** (1.9)	64%	*** (1.9)	15%	** (0.5)	7%		7%		7%	44%
password-fill	264	58%		20%	*** (0.2)	11%	*** (0.4)	17%	** (1.9)	9%		43%	69%
control-negative	507	55%		29%	*** (0.3)	55%	*** (6.3)	4%	*** (0.4)	6%		5%	65%
red-interstitial	263	54%		54%		6%	*** (0.2)	16%		13%		12%	32%
payment-fill	264	52%		28%	*** (0.3)	24%		10%		10%		27%	60%
control-topchrome	**505**	**51%**	–	**51%**	–	**23%**	–	**10%**	–	**9%**	–	**6%**	**67%**
install-content	264	45%		24%	*** (0.2)	45%	*** (3.7)	7%		19%	*** (8.7)	5%	27%
permission-quiet	264	44%	* (0.7)	44%		10%	*** (0.3)	22%	*** (2.6)	14%	** (3.2)	10%	44%
autofill-fill	263	43%	* (0.7)	43%	* (0.7)	9%	*** (0.3)	29%	*** (4)	10%		9%	63%
autofill-save	264	43%	* (0.7)	43%	* (0.7)	19%		19%	** (2.1)	11%		8%	52%
filepicker	506	39%	*** (0.6)	12%	*** (0.1)	27%		39%	*** (6.6)	14%	** (2.6)	8%	51%
fedcm-login-other	265	33%	*** (0.4)	33%	*** (0.4)	24%		14%		14%	** (2.9)	15%	50%
permission-context	264	28%	*** (0.3)	28%	*** (0.3)	35%	*** (2.1)	22%	*** (2.6)	8%		6%	68%
permission-plain	264	27%	*** (0.3)	27%	*** (0.3)	41%	*** (3.2)	19%	** (2)	7%		6%	71%
permission-legacy	264	26%	*** (0.3)	26%	*** (0.3)	48%	*** (4.3)	18%	** (1.9)	5%		3%	79%
install	264	20%	*** (0.2)	20%	*** (0.2)	49%	*** (4.4)	10%		15%	*** (4)	6%	33%
permission-os	263	16%	*** (0.1)	59%	* (1.4)	11%	*** (0.4)	16%	* (1.8)	8%		6%	73%
Overall	**2,024**	**53%**		**37%**		**21%**		**14%**		**9%**		**19%**	**59%**

C.3 Model Overview

Table 8: Impact of the five attributes among which the tested surfaces varied in study 1 on correct attribution. Two generalized linear models were run separately for mobile and desktop surfaces. Correct attribution was used as the dependent variable, while the five attributes listed in the table served were entered as independent variables. Respondent identifiers were entered as a random effect. The 'sig' columns indicate if this proportion is statistically significantly different (*< .05, **< .01, *< .001) from the respective reference level. Reference levels: Familiarity – Not at all, slightly, or somewhat familiar; Brand cues – none; Contains info – none; Active trigger – false; Overlap – none.**

		Desktop		Mobile	
Attribute	**Value**	odds ratio	sig.	odds ratio	sig
(intercept)		0.71	**	0.28	***
Familiarity	very or extremely familiar	2.09	***	2.06	***
Brand cues	Chrome logo	1.38	***	1.25	
	Google logo	2.27	***	6.51	***
	illustration	1.73	***	6.30	***
Contains info	user data	1.09		1.09	
	website	0.34	***	0.64	***
	both	0.92		0.92	
Active trigger	true	0.84	*	1.71	***
Overlap	centered	1.18		–	–
	low overlap	1.67	***	2.36	***
	high overlap	1.51	***	3.29	***

D Study 2: Questionnaire

The following questions were used to screen participants for eligibility:

Introduction Welcome to our survey! Thank you for taking the time to share your responses today.

This survey should take about 10-15 minutes, and we would appreciate your full attention.

We recommend taking the survey on a larger screen (laptop, tablet, etc.) so you're able to clearly see the images and have the best and quickest survey experience.

Q1. Do you work as a consultant or are you affiliated in some way with any of the following organizations or types of organizations?

- An advertising agency (Reject)
- A market-research firm or marketing department (Reject)
- A company that manufactures, distributes, or sells electronics (Reject)
- A government agency (Reject)
- None of the above (Accept)

Q2. What kind of mobile phone do you spend the most time using?

- Android (Accept)
- iOS (Accept)
- Not sure (Reject)

Q3. What model type is your phone?

- Samsung (Accept)
- Pixel (Accept)
- Motorola (Accept)
- OnePlus (Accept)
- LG (Accept)
- iPhone (Accept)
- Other (free response) (Reject)

Q4. Which age group best describes you?

- Under 18 (Reject)
- 18-23 (Accept)
- 24-30 (Accept)
- 31-40 (Accept)
- 41-50 (Accept)
- 51-60 (Accept)
- 60+ (Accept)
- Prefer not to answer (Reject)

Q5. Which of the following statements best describes your views on privacy online?

- I feel that companies should not be able to acquire my personal information for their own needs... (20% quota)
- I weigh the potential pros and cons of sharing my information... (65% quota)
- I am not very concerned with privacy... (15% quota)

Q6. Overall, how would you describe your knowledge of technology related to the Internet, computers, software, smartphones, and tablets?

- I am very knowledgeable about technology related to the Internet... (75% quota)
- I am not very knowledgeable about technology... (25% quota)

Q7. Imagine the following scenario: You download a new shopping app. So, you open the app for the first time and get prompted with the following:
[stimuli of permission prompt]
If you selected an option and then wanted to change your selection at another point in time, what would you be likely to do first?

- Open the app and look for settings within the app
- Open my phone's settings and look for the settings for this specific app

[next page]

Q8. Which entity wants to access your microphone?

- Android
- My device (e.g., Google, Samsung, OnePlus, etc.)
- The app
- Other (written response)
- I'm not sure

Q9. Which entity put this permission request pop up on your screen?

- Android
- My device (e.g., Google, Samsung, OnePlus, etc.)
- The app
- Other (written response)
- I'm not sure

[next page]

Q10. How likely do you think it is that Android wants access to your microphone?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q11. How likely do you think it is that your device (e.g., Google Pixel, Samsung, OnePlus, etc.)wants access to your microphone?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q12. How likely do you think it is that the app wants access to your microphone?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q13. How likely do you think it is that Android put this permission request pop up on your screen?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q14. How likely do you think it is that your device (e.g., Google Pixel, Samsung, OnePlus, etc.) put this permission request pop up on your screen?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q15. How likely do you think it is that the app put this permission request pop up on your screen?

- Not at all likely
- Slightly likely
- Somewhat likely
- Fairly likely
- Very likely

Q16. How confident are you in your prior responses?

- Not at all confident
- Slightly confident
- Somewhat confident
- Fairly confident

E Study 2: Additional Stimuli

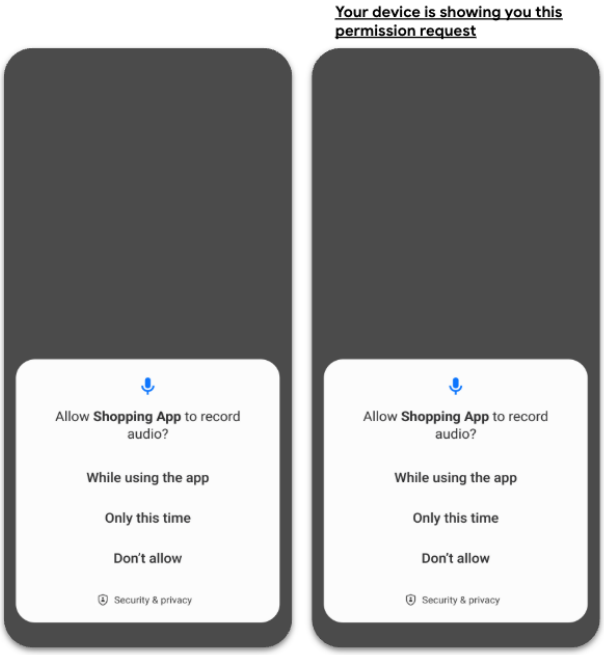

Figure 19: The image above shows the two S&P conditions for the Samsung group, which we will refer to as "S&P branding" and "S&P branding + explanatory text".