Websites Need Your Permission Too – User Sentiment and Decision-Making on Web Permission Prompts in Desktop Chrome

Marian Harbach
mharbach@google.com
Google
Munich, Germany

ABSTRACT

The web utilizes permission prompts to moderate access to certain capabilities. We present the first investigation of user behavior and sentiment of this security and privacy measure on the web, using 28 days of telemetry data from more than 100M Chrome installations on desktop platforms and experience sampling responses from 25,706 Chrome users. Based on this data, we find that ignoring and dismissing permission prompts are most common for geolocation and notifications. Permission prompts are perceived as more annoying and interrupting when they are not allowed, and most respondents cite a rational reason for the decision they took. Our data also supports that the perceived availability of contextual information from the requesting website is associated with allowing access to a requested capability. More usable permission controls could facilitate adoption of best practices that address several of the identified challenges; and ultimately could lead to better user experiences and a safer web.

CCS CONCEPTS

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

ACM Reference Format:
Marian Harbach. 2024. Websites Need Your Permission Too – User Sentiment and Decision-Making on Web Permission Prompts in Desktop Chrome. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3613904.3642252

1 INTRODUCTION

The web as a platform for using services and accessing information is becoming increasingly powerful [8]. Just like on other common application platforms, such as Android, iOS, Windows, or macOS, adding new capabilities often carries risks which cannot be entirely mitigated using purely technical means. In many cases, users are then asked if a given website or application should be allowed to access such a capability using permission prompts, turning this into a usable privacy and security problem. Permission prompts on mobile operating systems have been subject to usable security and privacy research for many years [6, 7, 10, 12, 17]. However, to the best of our knowledge, no prior work has investigated how permissions are used and understood by users on the web in general. This work aims to fill that gap and provide an in-depth overview of users' behavior, decision-making, and sentiment on web permission prompts. We identify opportunities for browser vendors and permission system designers to create more usable controls that are easier to integrate for web developers and better tailored to user needs while browsing the web, which can lead to less annoyance and better decision making, and ultimately help users to stay safe.

While many applications are offered on both the web and as apps on mobile platforms, the web enables use cases that are more ephemeral than those enabled by apps. Searching and information gathering are popular activities on the web, during which users will often interact with many websites with distinct owners. In contrast, before a user can use a mobile app, they need to choose it in a store and install it on their device, often confirming their choice with an authentication step. Intuitively, these actions alone seem to be correlated with permanence and trust. Apps are also often built to enable more permanent use cases, saving settings, log-in information and other preferences from the get-go. On the web, such permanent use cases also exist (e.g., with web-based email clients or productivity apps). Yet many websites are also used only very briefly.

Another important difference lies in the delivery of software [1]: websites are delivered dynamically and there's no guarantee that any two users will receive the same version of a website's source code. With apps, developers usually create bundles that are then delivered to all users equally until an update becomes available. Furthermore, there is no requirement to access websites in a certain way: users are free to choose their browser or default search engine, and to type any URL into the address bar of their browser. With apps, distribution happens via app stores, where certain requirements and guidelines are enforced. For permissions in particular, both Android [13] and iOS [3] specify minimum requirements and best practices that get enforced during the app store review process. No such lower bound exists for the permission user experiences that web developers can create. Using Javascript APIs, developers are free to show a permission prompt for a desired capability at any time during their website's control flow, even when a page is still loading. From the authors' own experience of the web, it seems apparent that there is less adherence to best practices, such as showing a permission prompt only after a relevant user interaction.

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05.
https://doi.org/10.1145/3613904.3642252

To mitigate some of these challenges, browser vendors have implemented various mechanisms to somewhat limit websites' abilities to show permission prompts. For example, Mozilla Firefox only allows asking for notification permission after the user interacted with the current page within 5 seconds [19]. The underlying assumption is that after such an interaction, it is more likely that the permission prompt will be relevant and associated to a user's intent. As another example, Chrome intervenes on websites' requests to show notification and geolocation prompts, as introduced by Bilogrevic et al. [4, 14].

Finally, the prompt UIs used in web browsers, especially on desktop platforms, are also quite different from their mobile counterparts. Mobile permission prompts are modal and very prominent on the screen, due to the often limited screen size. They also only show buttons with options to grant or deny access. In contrast, many desktop browsers including Chrome, Firefox, and Edge, show a lightweight, non-modal permission prompt anchored to the address bar. Additionally, there is an affordance to dismiss permission prompts (for example, by clicking an "x" button) and the prompt can be ignored entirely by just navigating away. To the best of our knowledge, these additional user actions on permission prompts have not been investigated before.

The first goal of this work is thus to describe and quantify the experience of permission prompts on the web. On the one hand, we describe current user behavior when encountering prompts. On the other hand, we aim to measure how annoyed and interrupted users feel, given that best practices are not as enforceable on the web as they are on mobile platforms [3, 13].

Inspired by the work of Bonné et al. and Cao et al. [6, 7], who looked at permission prompts on Android, a second goal of this work is to understand the reasons users of web permissions cite for making their decisions across all four possible outcomes (accept, deny, dismiss, and ignore). Assuming that permission systems aim to encourage "good" decisions, we delineate decision behaviors that appear rational (and thus should be encouraged) from behaviors that appear problematic (and the permission experience may have to be changed to avoid them).

A third goal is to investigate to what extent known factors impact users' decision-making on web permission prompts. Privacy theories such as contextual integrity [20] suggest that contextual information plays an important role, and prior work on mobile permissions [6, 10, 27, 31] found this to hold in practice. In particular, why a capability is requested and to what extent it is tied to the user's need in the current context influences users' decisions and perceptions. Prior work also argued that tying permission granting to intentional user interactions improves the availability of contextual information and thus helps users make better decisions [18]. We thus also explore to what extent the prior user interaction heuristic employed by Firefox is associated with a better user experience and better decision-making on the web today.

In sum, our work addresses the following research questions:

RQ1. How annoying or interrupting is prompting for permission on the web?
RQ2. How easy or difficult do users find making decisions on web permission prompts?
RQ3. Which reasons do users cite for their decisions? To what extent can we consider users' current decision-making rational?
RQ4. How much contextual information do users currently perceive prior to seeing permission prompts? Is more contextual information associated with rational decision-making?
RQ5. Does a perceived self-benefit from a capability impact the action users take on permission prompts?
RQ6. Does user interaction prior to a permission prompt impact any of the above?

As a first step towards understanding web permissions, the work presented in this paper focuses on the experience of web permissions on desktop browsers. We believe this is a valuable contribution, as the user interfaces of permission prompts on desktop platforms are most different from permission prompts on mobile operating systems.

The key contributions of this work comprise:

- We analyzed telemetry from more than 100M Chrome installations and find that ignoring and dismissing permission prompts are the most common actions Chrome users take. Our data supports that many websites are not adhering to best practices available for other platforms [3, 13], as 83.9% of permission prompts for the four most common capabilities shown to users are not preceded by a user interaction on the given website within 5 seconds.
- We collected responses from 25,706 Chrome users immediately after they made a decision on permission prompts using an experience sampling approach. Prompts are more likely to be annoying for geolocation and notifications requests, especially when requests are ultimately denied or dismissed. This implies that prompts where users do not find capability access necessary are more problematic.
- We provide an overview of reasons respondents cite for making decisions across four permission types and four actions users can take on a permission prompt, and compare them to prior findings on Android. We find that a majority of respondents cite at least one rational reason for their decision-making, even when dismissing and ignoring prompts.
- We investigate to what extent the perceived availability of contextual information or a perceived self-benefit is associated with actions users take on the four most common permission types. We find that respondents are more likely to have sufficient contextual information or perceive a self-benefit when microphone and camera permissions are requested as well as when they are allowing permission requests of any type.
- We investigate to what extent a prior user interaction within five seconds on a given website is associated with different behavior, improved sentiments, or rational decision-making. We find more users allow capability access after a prior user interaction in telemetry (a three-fold increase for geolocation prompts, for example) and observe a small effect of respondents being somewhat more likely to rate a prompt as not annoying in experience sampling responses. However,

we find no substantial effects on ease of decision-making, rational reasons for making a decision, certainty why the site is asking, or perceived benefit to oneself. We conclude that the prior user interaction heuristic can predict allowing behavior to a limited extent, but find no evidence of it being indicative of availability of contextual information and thus more rational decision-making.

The remainder of the paper is structured as follows: Section 2 provides additional details on how permissions work on the web today and how they compare to permission prompts on mobile platforms. Additionally, we provide an overview of prior research and how it relates to our work. Section 3 gives an overview of telemetry data, discussing which capabilities are commonly used on the web and which decisions users make for them. Section 4 then introduces our experience sampling methodology before Section 5 describes the findings. Section 6 lists limitations of the approach we chose before Section 7 provides a summary and discusses implications. Section 8 outlines next steps, and the Appendix provides the full questionnaires we used for the experience sampling questionnaires.

2 BACKGROUND AND RELATED WORK

In this section, we will provide a brief overview of permissions on the web and how they differ from permissions on mobile platforms, as well as related prior work on permissions in general.

2.1 Permissions on the Web

Over the past 15 years, the web platform has added capabilities and APIs to allow for increasingly powerful websites and web applications (jointly referred to as websites in this paper for brevity) [8]. Many of these capabilities need to access sensitive information outside of the browser sandbox and thus inherently carry a risk. There are technical mitigations, such as requiring a secure context to make sure network-based attackers cannot easily access a capability [28] or requiring transient user activation (see next paragraph). Beyond that, browsers rely on users to confirm that using the capability on a given website matches their intention. These confirmation questions are shown as permission prompts (see Figure 1 for examples), exposing a website’s request as a choice to users.

As described by Bilogrevic et al. [4], the push notification capability in particular has been so heavily used that it became bothersome for users. Chrome and other browsers started intervening on permission requests by websites to reduce the burden on its users. Notably, Firefox started requiring a prior user interaction on a given website within five seconds in version 72 before a permission prompt for notifications can be shown [19]. The overeager use of notifications illustrates a challenge that exists because of the open nature of the web. In contrast to other platforms where users commonly encounter permissions, like Android or iOS, there is no review process or gatekeeper that can enforce certain best practices or guidelines on the web. There also is no direct recourse for misbehaving websites, beyond mitigations such as Google Safe Browsing.

Additionally, there are some practical differences between permission prompt UIs on mobile devices and in browsers for desktop platforms. Prompts on mobile platforms offer additional controls,

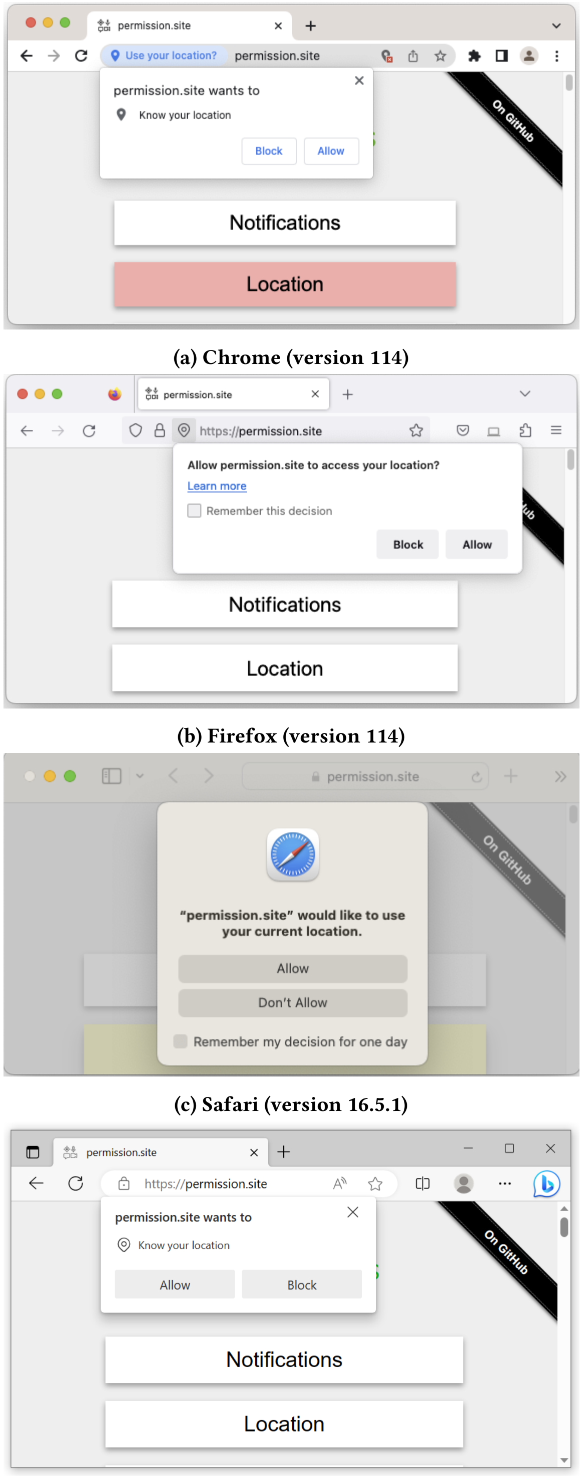

permission.site wants to

Know your location

(a) Chrome (version 114)

Allow permission.site to access your location?

Learn more

□ Remember this decision

(b) Firefox (version 114)

“permission.site” would like to use your current location.

□ Remember my decision for one day

(c) Safari (version 16.5.1)

permission.site wants to

Know your location

(d) Edge (version 114)

Figure 1: Permission prompt for geolocation access in four common web browsers.

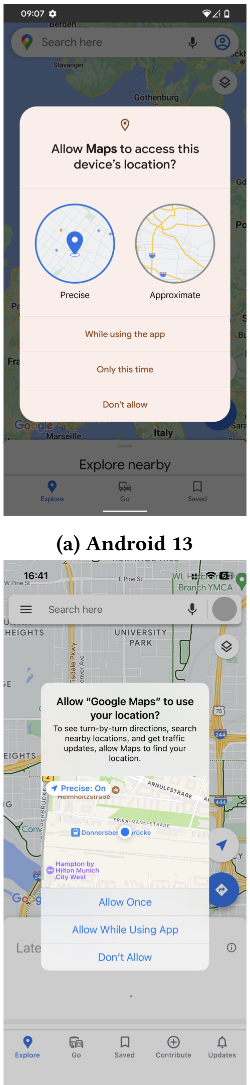
Allow **Maps** to access this device's location?

(a) Android 13

Allow “Google Maps” to use your location?

To see turn-by-turn directions, search nearby locations, and get traffic updates, allow Maps to find your location.

(b) iOS 16.5.1

Figure 2: Permission prompt for geolocation access on mobile platforms.

like coarse or fine location, developer-provided rationale strings and one-time allow options (see Figure 2b). Some desktop browsers also have one-time options (e.g., Firefox and Safari), but none offer developer-provided rationale strings or a coarse-fine toggle (see Figure 1). Mobile permission prompts are also modal. This means they appear in the middle of the screen, hiding the app's content until a decision has been made on the permission request. On iOS, one has to make an explicit allow or deny decision using the offered buttons. On Android, users can make the prompt disappear by tapping outside the prompt or by using the back button or gesture. These interactions do not have a visual affordance and, to the best of our knowledge, have not been investigated in prior work.

Permission prompts in many desktop browsers are presented quite differently: In Chrome, Edge, and Firefox, prompts hang off the address bar. Given the often larger screen sizes on desktop platforms, they can easily be ignored by just continuing what one was doing in the content area. This is referred to as the “ignore” action in the remainder of the paper. Additionally, Chrome, Edge, and Firefox afford a temporary block decision (clicking the “x” button in Chrome and Edge, and clicking the Block button without checking “remember my decision” in Firefox). We refer to this action as “dismiss” in the remainder of the paper. For both, ignore and dismiss decisions, the permission-gated capability remains inaccessible for the user's current visit, but the website gets to ask for permission again on the next visit.

Users may also choose to deny the request, in which case we again see differences between desktop browsers. Chrome's deny decisions are always permanent and thus the capability remains inaccessible not only for the current visit, but also for future visits, as the website will not be able to ask again. Firefox offers a “Remember my decision” checkbox to make decisions permanent, while they are temporary by default. Safari goes even a step further: for example, any permission decision on geolocation access can at most be remembered for one day. After a temporary decision expires, the site can ask again. For the remainder of this paper, we focus on Chrome, as a popular desktop browser at the time of writing [24].

2.2 Prior Research on Permission Prompts

A sizeable body of work has investigated permission prompts, primarily for mobile apps in general and on the Android platform in particular. Initially, Android relied on an install-time permission model, asking users to agree that an app gets to use all capabilities it desires as a condition of installing it. Numerous authors described the problems users have with this model [11, 12, 17].

With Android 6.0, the platform switched to a runtime permission model. Bonné et al. [6] investigated user behaviors and decision-making of Android runtime permissions, identifying differences in deny rates across permission types as well as a number of reasons underlying participants' decision-making. They find that an application's need for a given capability was a main reason for granting or denying, while being able to change the decision later was a common reason for denying. These findings were largely confirmed by Cao et al. [7] in 2021 with an experience sampling approach on users recruited via online advertising.

Beyond describing user behaviors, several authors have argued for contextualizing requests for access to capabilities or resources. Work by Thompson et al. [26] showed that indicators can help users to understand when access occurs and thus help to make sure access is appropriate. As a theoretical foundation, Nissenbaum's Contextual Integrity framework [20] postulates that privacy manifests as appropriate information flows, where an information flow is characterized by 5 properties (sender, receiver, data subject, data type, and transmission principles). Contextual norms (including people's expectations and social norms) are applied to these five properties and result in an appropriateness judgement. The implication is that if any of these five properties change, the information flow may no longer be appropriate. A corollary is that one needs to be aware of these five properties of an information flow to make informed judgements of appropriateness. Votipka et al. [27] confirm that when, why and with whom data is shared mattered when

Table 1: Chrome telemetry from desktop platforms on the most commonly used capabilities and user actions. % Prompts are calculated against the total number of all prompt events across all capabilities. One exception are the two percentages in parentheses in the Total row, denoting the percentage of prompts with and without user interaction for the four capabilities in the table only. Prompts per 1k page loads describes the frequency with which the given prompt type is shown. User action rates are computed row-wise, i.e. as percent of number of prompts for a (capability, prior user interaction) pair.

	Prior User	% Prompts	Prompts per	User Action			
Capability	**Interaction**	**of overall total**	**1k page loads**	ignored	dismissed	allowed	denied
Notification	no	43.0%	1.9	42.9%	36.6%	10.4%	10.2%
	yes	8.0%	0.3	29.9%	37.5%	19.9%	12.7%
Geolocation	no	31.2%	1.3	53.8%	30.9%	9.0%	6.3%
	yes	4.4%	0.2	28.3%	34.5%	27.4%	9.8%
Microphone	no	2.2%	0.1	9.2%	24.8%	60.9%	5.1%
	yes	1.6%	0.1	3.9%	24.8%	65.1%	6.2%
Camera	no	0.6%	0.03	9.2%	18.5%	66.8%	5.6%
	yes	0.8%	0.03	2.8%	20.9%	70.3%	5.9%
Total	no	77.0% (83.9%)	3.3	46.1%	33.8%	11.7%	8.4%
	yes	14.8% (16.1%)	0.6	25.1%	34.3%	29.8%	10.8%
	overall	91.8%	4.0	42.7%	33.9%	14.6%	8.8%

making access decisions in the context of mobile apps. Wijesekera et al. [31] find that at least 80% of participants wanted to prevent at least one access to data given a specific context and that the use of an ML-based classifier to contextualize and automate permission decisions reduces the number of unexpected capability accesses by 75% [32]. More recently, Elbitar et al. [10] found that providing rationales with permission requests appears beneficial for both users and developers. Cao et al. [7] find that the presence of a rationale string halved the deny rate of permissions requested by the apps in their sample. They also find that deny rates are lower when their participants expected an app to ask for the given permission.

Additional work looked at mobile app permissions from a developer perspective and suggested user interaction improvements. Tahei et al. [25] interview developers and compare their views with end users' mental models. Harbach et al. [15] propose to use concrete examples of information being made available to an app to help with assessing risks associated with permission grants. Micinski et al. [18] propose to tie interactive usage of capabilities to user interactions, finding that capability access is more expected after an explicit user interaction.

Very little work has addressed web permissions directly. In 2012, Chaitrali and Traynor [1] discussed differences in mobile web applications and native apps on mobile operating systems, citing the dynamic nature of the delivery of websites, the absence of app stores, and the ease with which users can happen upon a random website, among others, as core differences between web and native applications. They propose to improve inspection abilities via a manifest-type mechanism and call for alignment with operating-system level permission systems to help users make informed decisions. Hazhirpasand et al. [16] present a click-jacking attack on web permission prompts. Bilogrevic et al. [4] and Harbach et al. [14] built and evaluated an intervention to reduce user interruptions due to overeager websites prompting for permissions frequently for both mobile and desktop Chrome.

To the best of our knowledge, no prior work has investigated user sentiment and behavior for permissions on the web in general and on desktop platforms in particular.

3 TELEMETRY OF PERMISSIONS ON THE WEB

We use telemetry provided by Chrome installations on desktop platforms to outline the status quo of user behavior on web permission prompts. Telemetry data is available when users did not opt out of the collection via the "Help improve Chrome's features and performance" toggle in Chrome settings or during Chrome's installation or first-run flow. We exclude any prompts with interventions (when Chrome proactively hides the prompt because of a user preference, because Chrome determined users are unlikely to allow, or because users repeatedly dismissed or ignored prior prompts; see Harbach et al. [14]), as they were also excluded when showing experience sampling prompts (see Section 4). As a first step, we focus on desktop platforms, given how different the user experience of permission prompts is there. It is noteworthy that use cases when browsing the web can be different between desktop and mobile devices [22] and understanding the situation on mobile devices is thus important future work.

We report data aggregated over a 28-day window ending on Jun 30, 2023. The data set contains data from hundreds of millions of permission prompts from more than 100 million Chrome installations. *Prior user interaction* refers to a transient activation event happening before the permission prompt was shown, i.e. the user had a mouse click anywhere on the content area or a keyboard event within 5 seconds before the prompt appeared [30]. We also report four user actions as outcomes of the prompt being shown: *allowed* and *denied* prompts are those where the user clicked on the respective button in the prompt (see Figure 1a); *dismissed* prompts are those where the user clicked the "x" icon at the top right of the prompt; and *ignored* prompts are those where the user did nothing and navigated away or closed the tab or browser window.

Table 1 provides a breakdown of permission prompts across the capabilities that triggered the respective prompt. Notifications, geolocation, microphone, and camera access account for 92% of all shown permission prompts and we thus focus our analysis in this section and the remainder of the paper on those capabilities.

Looking at notifications and geolocation in one group and microphone and camera in a second, we can see that user behaviors

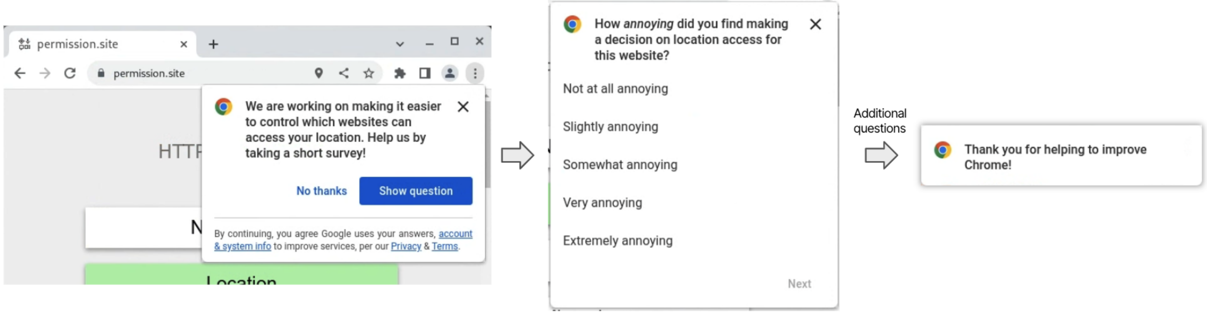

We are working on making it easier to control which websites can access your location. Help us by taking a short survey!

By continuing, you agree Google uses your answers, account & system info to improve services, per our Privacy & Terms.

How *annoying* did you find making a decision on location access for this website?

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

Thank you for helping to improve Chrome!

Figure 3: Screenshot of questionnaire invitation and subsequent screens. The final ’Thank you’ message disappeared automatically after 5 seconds.

between those groups are different. Whereas the former shows higher rates of ignoring and dismissing, the latter group is most commonly allowed. This difference in itself suggests that the use cases associated with the respective capabilities lead to different user behavior. We can speculate that the nature of the capabilities in question make them more or less central to the underlying use cases and thus result in differences in users’ perceived need of the capability. In addition, across all capabilities, a prior user interaction leads to lower ignore and higher allow rates. These increased allow rates suggest that differences in the permission request experience lead to different decision behaviors.

It is also noteworthy that only 16.1% of prompts for the four most common capabilities were shown after a user interacted with the page within in 5 seconds in the 28-day window covered by our data. In these cases, ignore rates are 21.0% lower and allow rates 18.1% higher overall. Allow rates for the geolocation capability stand out in particular, increasing three-fold from 9.0% without to 27.4% with a prior user interaction.

As a point of comparison, prior work on Android [6, 7] found deny rates of about 14% for the location permission, 14-16% for camera and 26-30% for microphone in their samples, which are substantially higher than what we find on the web. Given that desktop browsers allow other non-allow outcomes, this is not surprising. As Android apps could show notifications by default until Android 13 [2], notification prompts were not investigated in prior work on Android.

4 EXPERIENCE SAMPLING METHOD

Based on the user behavior observed using the telemetry outlined in Section 3, we set out to answer the research questions introduced in Section 1. Given the short-lived and contextual nature of permission decision moments and the research approaches taken in prior work, we chose to use experience sampling to gather data from Chrome’s users.

Chrome users are eligible to see an experience sampling prompt when all of the following conditions for their Chrome profile are met:

- Not opted out of “Help improve Chrome’s features and performance” setting;
- Not displayed another experience sampling prompt within 180 days;
- Created profile or installed Chrome at least 30 days ago;
- Chrome is not recovering from a crash; and
- A questionnaire language matching the current Chrome language (locale) is available.

The invitation to the questionnaire showed approximately five seconds after a user made a choice on a permission prompt (and the prompt thus disappeared). This delay is due to a technical limitation where questionnaires cannot be pre-fetched because of server load constraints. If the conditions for showing a questionnaire were met, users would first see an invitation page, and then one question per page (see Figure 3). Respondents were able to abandon the questionnaire at any time by clicking the "x" button in the top-right corner.

4.1 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 Google 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 retain any identifying data with our questionnaires. Participation in experience sampling studies in Chrome is only offered to users at most once per 180 days and only if they did not opt out of sharing telemetry data. Each questionnaire we fielded was short and easy to ignore or dismiss. The questionnaire invitation provided links to Google’s privacy policy as well as an overview of any additional data sent along with their responses. This data comprised which permission type they saw a prompt for, the action they took on it, the type of prompt UI they saw, whether there was a user interaction prior to the prompt showing (as defined in Section 3), their user agent string, the current timestamp, and their timezone offset.

4.2 Statistical Testing

To compare response proportions between various slices of the data, we use omnibus χ^2 tests and report pairwise differences when the absolute value of standardized residuals (*sresid*) is at least two [23]. We also use logistic regressions to test for differences caused by several nominal variables on binary outcomes. Independent variables used comprise the requested capability (reference: “notifications”),

Table 2: Overview of response behaviors from the first questionnaire. Accepted questionnaires are those where respondents answered at least one question. They are counted as partial if not all questions were answered, as complete otherwise. The rightmost column shows the fraction of accepted questionnaires that followed a permission prompt being shown after a user interaction.

User action	Capability	Invitations shown	% accepted	# partial	% partial	# complete	Median time [sec]	% w/ user interaction
allowed	notifications	45,201	2.8%	282	22.1%	995	40	43.6%
	geolocation	33,770	3.5%	160	13.5%	1,025	38	39.3%
	microphone	49,077	2.9%	375	26.4%	1,044	38	50.6%
	camera	47,294	2.9%	337	24.4%	1,044	38	56.7%
denied	notifications	50,249	2.3%	125	10.7%	1,038	42	23.9%
	geolocation	55,437	2.1%	136	11.5%	1,043	41	21.6%
	microphone	16,664	2.7%	105	23.6%	340	34	61.1%
	camera	9,584	2.4%	51	22.6%	175	31	62.3%
dismissed	notifications	165,204	0.8%	240	18.8%	1,037	46	21.7%
	geolocation	148,741	0.7%	179	16.3%	921	48	18.9%
	microphone	52,287	1.0%	138	26.7%	379	39	54.3%
	camera	22,505	1.0%	59	27.2%	158	37	59.1%
Total		696,013	1.6%	2,187	19.2%	9,199	40	40.3%

Table 3: Overview of reasons offered to respondents when asking about decision reasons. "2x" indicates that two offered items matched the reason in these conditions.

		Included for user action			
Category	**Reason**	allow	dismiss	deny	ignore
rational	want to allow/decide later		x	x	x
	functionality	x	2x	2x	2x
	developer trust	x	x	x	x
so-so	nothing bad will happen	x			
	can't remember what I did	x	x	x	x
problematic	did not notice				x
	won't work otherwise	x			
	want the popup to go away	x	x	x	x
	won't be able to allow later	x			

the action the respondent took (reference: "allowed") as well as the presence of a user interaction before the prompt showed (reference: no user interaction). The independent variables we include were selected based on our research questions. We did not optimize these exploratory models any further. Result tables for the regressions include a "sig." column, that indicates statistical significance levels, with . = $p < .1$, * = $p < .05$, ** = $p < .01$, and *** = $p < .001$.

4.3 Fielding & Responses

We launched two distinct campaigns to address RQs 1-2 and RQs 3-6 separately, as we wanted to keep the individual questionnaires short. Both questionnaires were fielded to Chrome installations on Windows, macOS, ChromeOS, and Linux using an English language setting. We aimed to collect approximately 1,000 responses for each (capability, user action) pair to retain sufficiently large subgroups when slicing the data during analysis in both questionnaires. In both questionnaires, the surveyed capabilities comprised the four most used capabilities on the web (notifications, geolocation, camera, and microphone; see Section 3). The user actions included in each questionnaire are described below. The analysis was based on complete responses; partial responses were discarded.

4.3.1 Questionnaire 1. As RQs 1 and 2 are about user sentiment during interaction with the prompt, we showed the questionnaire to users who actively interacted with the prompt, i.e. dismissed, allowed, or denied it but not those who ignored it. The questionnaire thus had a 4x3 between subjects design. We used three simple Likert-scale-type rating questions (see Appendix A.1) to measure annoyance, interruption, and ease of use.

Questionnaire 1 was active between August 31st and November 17th, 2022. We collected a total of 9,199 complete responses. Median questionnaire completion time was 40 seconds and Table 2 provides an overview of response behaviors. Due to the lower number of prompts on microphone and camera prompts (see Section 3), lower questionnaire accept rates on dismissed prompts, as well as higher fractions of partial responses on microphone and camera questionnaires, we were unable to meet the desired response counts in the available time for several conditions. This behavior suggests that some capabilities and user actions occur in situations where users are less amenable to completing a questionnaire. Overall, the accept rates on our in-product questionnaires are similar to other experience sampling questionnaires Chrome shows on unrelated features.

4.3.2 Questionnaire 2. After the challenges with gathering sufficient responses during questionnaire 1, we increased the size of the population that would see conditions that filled too slowly. For this questionnaire, we included the "ignore" user action, as this questionnaire aimed to understand user decision-making across all outcomes. This questionnaire thus had a 4x4 between subjects design and comprised three questions (see Appendix A.2).

The first question asked participants to select which of a list of pre-defined reasons describe why they chose to take the action they took (RQ3). This multi-select question was modelled after the list of reasons used by Bonné et al. [6]. However, we had to substantially shorten their list to avoid overwhelming our participants, as they did not specifically sign up for this questionnaire and saw it in a small popup window on top of their regular browser window. We also adapted the list to match how permissions work on the web. As in the study of Bonné et al., the pre-defined reasons varied

Table 4: Overview of response behaviors from the second questionnaire. Columns are defined as in Table 2.

User action	Capability	Invitations shown	% accepted	# partial	% partial	# complete	# retained	Median time [sec]	% w/ user interaction
accepted	notifications	49,000	2.9%	240	17.2%	1,159	1,065	55	50.4%
	geolocation	37,860	3.5%	140	10.4%	1,201	1,152	51	39.0%
	microphone	39,423	3.8%	282	19.1%	1,198	1,115	49	49.9%
	camera	40,635	3.6%	258	17.7%	1,198	1,102	48	59.7%
denied	notifications	81,939	1.6%	134	10.0%	1,212	1,153	63	21.3%
	geolocation	89,807	1.5%	141	10.4%	1,211	1,168	59	19.2%
	microphone	32,607	2.2%	183	25.8%	525	472	46	55.3%
	camera	43,987	2.3%	224	21.7%	809	737	46	62.8%
dismissed	notifications	297,333	0.6%	291	17.1%	1,408	1,314	62	21.2%
	geolocation	281,130	0.6%	286	16.7%	1,425	1,355	58	17.7%
	microphone	101,683	0.7%	186	24.5%	572	534	46	56.6%
	camera	108,095	1.1%	272	23.8%	871	796	47.5	58.8%
ignored	notifications	292,065	0.4%	180	14.7%	1,043	942	63	17.2%
	geolocation	271,356	0.8%	303	13.7%	1,914	1,782	63.5	13.2%
	microphone	89,314	1.5%	293	21.3%	1,081	996	55	29.6%
	camera	54,455	2.0%	198	18.2%	888	824	53	24.2%
Total		1,910,689	1.1%	3,611	16.9%	17,715	16,507	55	33.8%

based on which action the respondent had taken and not all reasons were applicable to all user actions. For example, "nothing bad will happen" does not apply when not allowing, as there is no risk from exposing the capability in these cases. Table 3 provides an overview which reasons we asked about for which user action. The exact wording can be found in Appendix A.2.

Additionally, permission systems should aim to improve users' decision-making by rooting decisions in an "optimizing" response, according to Böhme and Grossklags [5]. Such decisions will be depending on information and judgement as opposed to reactions and heuristics. In contrast, finding that users cite reasons suggesting satisficing behavior would indicate "low motivation", "high difficulty of the question" or "monotonous repetition", all of which are problematic for permission systems.

Based on the optimizing vs. satisficing framework, and in an attempt to capture the status quo, we assign each reason to a category, labelling reasons that refer to information of the current context (functionality, trust in the developer) or an admittance of not having enough information at this point (wanting to allow/decide later) as "rational", denoting behavior of the optimizing kind. At the opposite end, we label reasons that show a lack of understanding (won't be able to allow later, won't work otherwise), a use of heuristics (wanted the popup to go away) or a failure to grab the user's attention (did not notice) as problematic, denoting behaviors associated with satisficing. We emphasize that it is not respondents' decision-making that is problematic, but it is problematic that Chrome's current permission user experience gives rise to such decision making. Finally, we label "nothing bad will happen" as so-so, since this can be an appropriate, fact-based judgement on, for example, well-known sites, but can also be a fatalistic heuristic in other cases.

The second question was designed to capture the amount of contextual information respondents were able to gather before making the decision (RQ4), while the third aimed to capture who respondents anticipated would benefit from granting access (RQ5). We added the last question to capture respondents' interpretations of the available contextual information in terms of perceived utility.

Questionnaire 2 was active from January 11th to March 27th, 2023. We collected a total of 17,715 responses. The median questionnaire completion time was 55 seconds and Table 4 provides an overview of response behaviors. Even though we increased the size of the population that saw the respective questionnaire versions, we were again unable to fill all conditions to the desired response count. The lower accept rates as well as higher abandon rates are consistent with the first questionnaire, though. We removed 1,208 responses for quality concerns, because these respondents had selected reasons from all three reason categories. The fraction of responses completed after a prompt with prior user interaction varies between capabilities and user actions, but is consistent with the first questionnaire.

5 FINDINGS

We will present findings for each research question (RQ) introduced in Section 1 below and discuss their implications in Section 7.

5.1 RQ1: Annoyance and Interruption

Overall, a majority of 63.9% of respondents did not find the permission prompt particularly annoying (either "not at all" or "slightly annoying"). Table 5 shows respondents' ratings across capabilities and user actions and Table 6 shows the results of a logistic regression on these factors. A detailed breakdown of respondents' rating of annoyance can be found in Figure 6 in Appendix B. Accepted prompts were rated as significantly less annoying. Additionally, microphone and camera prompts were also associated with significantly less reported annoyance. The former could be explained by accepted prompts being associated with a value that respondents desired. The latter may be related to geolocation and notification capabilities being substantially more commonly requested on the web and thus respondents seeing them more often. Camera and microphone prompts, on the other hand, may be associated with benefits that are immediate and easily understood by the user, such as being seen when joining a video conference; while notifications may only have a benefit at a later point in time, if at all.

Table 5: Responses to the questions on annoyance and ease of decision making as top-2-box scores ("not at all" or "slightly annoying"; "somewhat" or "very easy"). Two separate omnibus χ^2 tests were applied to each dependent variable. Blue and green cell background and * or ** indicate standardized residuals (sresid) > 2 and > 5 respectively. Yellow and orange cell background and § and §§ indicate sresid < -2 and < -5 respectively.

User Action	Capability	# not annoying	% not annoying	# easy	% easy
accepted	notifications	722	72.6%*	505	50.8%§
	geolocation	739	72.1%*	593	57.9%
	microphone	826	79.1%**	676	64.8%*
	camera	835	80.0%**	643	61.6%
denied	notifications	439	42.3%§§	651	62.7%
	geolocation	499	47.8%§§	699	67.0%*
	microphone	230	67.6%	196	57.6%
	camera	124	70.9%	102	58.3%
dismissed	notifications	578	55.7%§	486	46.9%§
	geolocation	522	56.7%§	506	54.9%
	microphone	251	66.2%	200	52.8%
	camera	112	70.9%	78	49.4%
Total		5,877	63.9%	5,335	58.0%
Omnibus χ^2			$\chi^2(11) = 671$, $p < .0001$		$\chi^2(11) = 156$, $p < .0001$

Table 6: Results of a logistic regression using feeling "not at all" or "slightly" annoyed as the dependent variable.

	Not annoying			
	Log odds	Std. Error	Odds	sig.
(Intercept)	0.78	0.054	2.19	***
Capability				
geolocation	0.10	0.053	1.10	.
microphone	0.55	0.069	1.73	***
camera	0.62	0.079	1.85	***
User Action				
denied	-0.99	0.056	0.37	***
dismissed	-0.63	0.057	0.53	***
Had prior user interaction	0.14	0.050	1.15	**

Furthermore, it is noteworthy that most denied notifications and geolocation prompts were still rated as annoying, even though Chrome's prompt quieting mechanism (as described by Harbach et al. [14]) had already made 41% and 19% of all notification and geolocation prompts respectively ineligible for this study: prompts for which Chrome determined that users were very unlikely to allow access and thus received a quieter UI treatment were excluded from our sample. However, even the remaining denied prompts still felt more annoying than other prompts to respondents.

We also asked a second, similar question about feeling interrupted instead of feeling annoyed in the first questionnaire. Responses to this question were strongly correlated with annoyance ratings (Spearman's $\rho = .74$, $p < .0001$) and exhibited the same patterns of significant differences. We thus don't report the values separately and conclude that feeling interrupted and annoyed are very similar sentiments when it comes to permission prompts.

5.2 RQ2: Ease of Decision-Making

Across all prompts we included in our study, 58% of respondents found it "somewhat" or "very easy" to make a decision on the permission prompt. Table 5 and the logistic regression results in Table 7 show that this was somewhat different across capabilities and user actions. A detailed breakdown of respondents' rating of ease of decision making can be found in Figure 7 in Appendix B. Respondents were 29%-37% more likely to rate geolocation, camera, and microphone decisions as easy. This may mean that it is more obvious when these capabilities are necessary or useful. Across user actions, denying was slightly more likely to be rated as easy while dismissed prompts slightly less.

Table 7: Results of a logistic regression using rating decision-making as "easy" or "very easy" as the dependent variable.

	Easy to make a decision			
	Log odds	Std. Error	Odds	sig.
(Intercept)	0.17	0.051	1.19	***
Capability				
geolocation	0.26	0.052	1.30	***
microphone	0.31	0.063	1.37	***
camera	0.26	0.070	1.29	***
User Action				
denied	0.23	0.054	1.26	***
dismissed	-0.28	0.054	0.75	***
Had prior user interaction	-0.04	0.046	0.96	

5.3 RQ3: Reasons for Decision-Making

In this section, we will look at the reasons respondents cited for their decisions. We split the analysis by user action, as the pre-defined reasons we provided to respondents also differed based on the action they took on the prompt (see Table 3). For accept and deny decisions, we compare our findings to those of Bonné et al. [6] based on their Android study conducted in 2016, whose findings were confirmed by Cao et al. [7] in 2021, finding very similar frequency of reasons for accepting and denying.

5.3.1 Reasons when Allowing. Table 8 shows that, when allowing a permission, functionality and developer trust reasons are most

Table 8: Responses to the question "When the website asked for access to $capability, why did you do what you did? Select all that apply." split by the capability requested by the website for respondents who eventually allowed access. As multiple selections were possible, percentages are based on number of respondents accepting the given capability and the "# respondents" rows show counts of unique respondents in each category.

Category	Reason	Notification	Geolocation	Microphone	Camera	Total
rational	developer trust	37.1% (395)	36.1% (416)	40.2% (448)	41.5% (457)	38.7% (1,716)
	functionality	29.8% (317)	51.8% (597)	52.6% (586)	51.0% (562)	46.5% (2,062)
	# respondents	*58.9% (627)*	*75.1% (865)*	*78.0% (870)*	*79.9% (881)*	73.1% (3,243)
so-so	can't remember what I did	15.8% (168)	5.2% (60)	7.6% (85)	4.6% (51)	8.2% (364)
	nothing bad will happen	13.3% (142)	11.9% (137)	14.4% (161)	14.4% (159)	13.5% (599)
	# respondents	*28.5% (304)*	*16.9% (195)*	*22.0% (245)*	*18.7% (206)*	21.4% (950)
problematic	wanted the popup to go away	11.2% (119)	7.3% (84)	2.9% (32)	3.7% (41)	6.2% (276)
	won't be able to allow later	4.0% (43)	2.1% (24)	3.4% (38)	3.2% (35)	3.2% (140)
	won't work otherwise	10.2% (109)	11.1% (128)	7.4% (83)	7.5% (83)	9.1% (403)
	# respondents	*22.5% (240)*	*18.7% (215)*	*12.4% (138)*	*13.7% (151)*	16.8% (744)
	other	4.0% (43)	3.6% (41)	0.8% (9)	1.8% (20)	2.5% (111)
Total	**# respondents**	1,065	1,152	1,115	1,102	4,434

Table 9: Overview of reasons provided by respondents who denied access on the corresponding permission prompt.

Category	Reason	Notification	Geolocation	Microphone	Camera	Total
rational	functionality	94.4% (1,088)	83.6% (976)	33.1% (156)	32.0% (236)	54.1% (1,910)
	want to allow/decide later	27.7% (319)	33.3% (389)	38.1% (180)	35.8% (264)	32.6% (1,150)
	developer trust	28.1% (324)	29.8% (348)	10.4% (49)	12.9% (95)	23.1% (816)
	# respondents	*88.8% (1,024)*	*89.3% (1,043)*	*62.3% (294)*	*63.1% (465)*	*80.1% (2,826)*
so-so	can't remember what I did	6.1% (70)	5.7% (67)	25.0% (118)	24.0% (177)	12.2% (432)
problematic	wanted the popup to go away	31.0% (357)	22.0% (257)	17.6% (83)	16.4% (121)	23.2% (818)
	other	4.0% (46)	4.5% (53)	7.6% (36)	8.4% (62)	4.5% (158)
Total	**# respondents**	1,153	1,168	472	737	3,530

common. 73.1% of respondents selected at least one rational reason when allowing. Notably, rational reasons related to functionality were selected less frequently when allowing notifications, which may be explained by the value of receiving notifications only manifesting at a later point in time. Additionally, in comparison to what prior work found on Android, functionality as a reason is generally less prevalent on the web: 68% selected "I want to use a specific feature that requires this permission" on Android [6] while 46.5% selected "I want to use a feature that requires $capability" in our study.

In the so-so category, the "nothing bad will happen" reason was selected by 13.5% of respondents, which is similar to what Bonné et al. found for allowing a permission on Android. Problematic reasons were selected by 16.8% of respondents. While low, 6.2% of respondents still just wanted the popup go away and thus allowed the permission (Bonné et al. found 10.2% on Android). Similarly, we find that "won't work otherwise" was selected by 9.1% of respondents, while Bonné et al. reported 23.8% for this reason.

In sum, it appears that rational decision-making dominates how respondents dealt with allowed permission prompts, which is very positive. Problematic reasons are still common, but apparently less so than on Android. We speculate that this may have to do with the availability of easy ignore and dismiss actions, thus deferring an answer to another time.

5.3.2 Reasons when Denying. Next, we will look at the same data but for respondents who denied the request for access to a capability they saw. As a reminder, in this condition, participants saw a similar but slightly different set of response options (see Table 3). Overall, we find that 80.1% of respondents selected a rational reason for denying. Table 9 shows that functionality reasons are substantially less common for denying microphone and camera access. At the same time, developer trust is much less common for those capabilities and more common for notifications and geolocation. Functionality-related reasons for denying a permission were also most commonly cited in Bonné et al.'s Android study (for example, 41% selected "I think the app shouldn't need this permission" [6]).

Respondents seeing requests for camera and microphone were more likely to state that they can't remember what they did. This may be explained by camera and microphone use cases possibly being more interactive and the permission prompt being less salient due to its placement and size, and thus easier to forget. The problematic "wanted the popup to go away" reason was more common for denied notification and geolocation prompts and in general much more commonly selected than when accepting a prompt. It is also more common in our study than it was in prior work on Android (13% selecting "I wanted the permission screen to go away" [6]). In line with the findings on annoyance in questionnaire 1 (see Section 5.1), this suggests that denying notifications is more annoying and respondents thus more frequently just want the popup to go away.

Table 10: Overview of reasons provided by respondents who eventually dismissed the corresponding permission prompt.

Category	Reason	Notification	Geolocation	Microphone	Camera	Total
rational	functionality	62.7% (850)	63.8% (838)	23.2% (124)	23.0% (183)	41.5% (1,658)
	want to allow/decide later	13.7% (186)	14.7% (193)	38.4% (205)	38.7% (308)	22.3% (891)
	developer trust	21.5% (291)	18.1% (238)	7.1% (38)	7.7% (61)	15.7% (628)
	# respondents	*70.8% (959)*	*69.7% (916)*	*59.4% (317)*	*61.9% (493)*	*67.1% (2,685)*
so-so	can't remember what I did	21.0% (285)	18.3% (241)	29.8% (159)	24.5% (195)	22.0% (880)
problematic	wanted the popup to go away	24.6% (334)	34.6% (455)	16.5% (88)	17.8% (142)	25.5% (1,019)
	other	4.0% (54)	2.9% (38)	8.6% (46)	9.0% (72)	4.4% (175)
Total	**# respondents**	1,355	1,314	534	796	3,999

Table 11: Overview of reasons provided by respondents who eventually ignored the corresponding permission prompt.

Category	Reason	Notification	Geolocation	Microphone	Camera	Total
rational	functionality	58.0% (546)	48.1% (858)	18.0% (179)	15.2% (125)	30.6% (1,390)
	want to allow/decide later	10.7% (101)	11.0% (196)	24.9% (248)	29.9% (246)	17.3% (788)
	developer trust	22.3% (210)	17.6% (313)	5.5% (55)	5.3% (44)	13.7% (622)
	# respondents	*62.0% (584)*	*53.5% (954)*	*41.6% (414)*	*46.4% (382)*	*51.4% (2,334)*
so-so	can't remember what I did	18.2% (171)	16.3% (290)	22.3% (222)	20.9% (172)	18.8% (855)
problematic	did not notice	23.1% (218)	36.1% (643)	31.3% (312)	25.4% (209)	30.4% (1,382)
	wanted the popup to go away	26.4% (249)	15.0% (268)	10.3% (103)	10.2% (84)	15.5% (704)
	# respondents	*44.3% (417)*	*47.8% (851)*	*41.1% (409)*	*35.0% (288)*	*43.2% (1,965)*
	other	2.2% (21)	5.2% (92)	11.4% (114)	15.8% (130)	5.9% (269)
Total	**# respondents**	942	1,782	996	824	4,544

5.3.3 Reasons when Dismissing. Dismissing is a common action on the web using desktop Chrome (see Section 3). To the best of our knowledge, reasons for taking this action have not been explored in any prior work for any platform. We find that the reasons respondents selected after dismissing are similar to when denying. Rational reasons are still most commonly selected, suggesting that dismissing is a helpful action respondents take intentionally. However, rational reasons are also somewhat less frequently selected in comparison to respondents that denied (67.1% vs. 80.1%). Especially for notification and geolocation requests, respondents selected reasons related to functionality substantially less frequently (63% and 64% vs. 94% and 84% when denying). For these two capabilities, we find that being unable to remember is more common after dismissing. Furthermore, we see similar levels of wanting the popup to go away, which is arguably less problematic when dismissing, as this decision is not permanent.

5.3.4 Reasons when Ignoring. Finally, on Chrome for desktop platforms, permission prompts can also be ignored by not interacting with them and then closing the tab or window or navigating away. We find that even when not taking explicit action, a majority of respondents indicated doing so with a rational reason related to functionality, wanting to decide later or a lack of trust in the developer of the website (see Table 11). However, the fraction of respondents citing at least one rational reason is lowest when compared to reasons for taking one of the other three actions on a permission prompt (51% vs. 67%, 80% and 73% when dismissing, denying, or allowing respectively). When ignoring, a possible, although problematic, reason for doing so is to not have noticed the prompt itself appearing. This reason was selected by 30.4% of respondents and is thus as common as ignoring because the requested capability is not related to desired functionality. This suggests that the lightweight UI that does not obstruct the content area (chosen by many desktop browsers with the notable exception of Safari) also has a substantial downside.

5.4 RQ4: Availability of Contextual Information

To investigate to what extent the perceived availability of contextual information plays a role in permission prompt decision-making, we asked participants how sure they were about why the current website asked for the given capability. Figure 4 shows that more respondents indicated that they were sure why the site asked when accepting a prompt within each capability. Across capabilities, it seems apparent that requests for camera and microphone offered more contextual information to participants, as more respondents indicated being sure across all four user actions. This is confirmed by a logistic regression as shown in Table 12, where respondents being asked for access to microphone or camera are 33% and 38% more likely to feel sure and those who decided to not allow access were 17-19% less likely to feel sure. Having had a prior user interaction only exhibits a very small effect on feeling sure why the site was asking for permission.

It seems at least plausible that those who felt sure why the website is asking would also be more likely to cite at least one rational reason for their decision. Running another logistic regression (see Table 13), we find such an effect, although not very strong: those who felt "very" or "extremely sure" why the website was asking were 7% more likely to cite at least one rational reason for their decision.

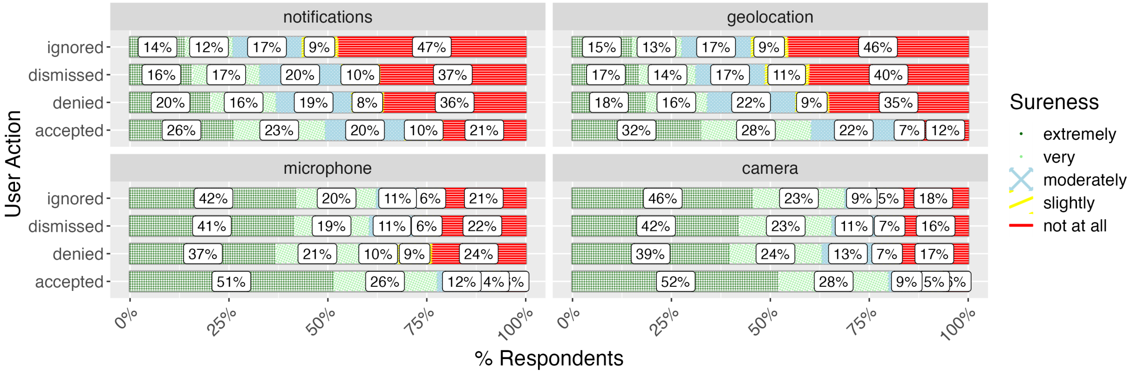

Figure 4: Respondents' ratings of how sure they are why the website asks for capability access (Q2 in questionnaire 2) across the requested capabilities and action they took on the permission prompt.

Table 12: Results of two logistic regressions using (1) feeling "very" or "extremely sure" why a website is asking for the given capability (Q2 in questionnaire 2) and (2) whether or not respondents believed there is benefit to themselves or to both themselves and the website for using the requested capability (Q3 in questionnaire 2) as dependent variables.

	Sureness				Self-benefit			
	Log odds	Std. Error	Odds	sig.	Log odds	Std. Error	Odds	sig.
(Intercept)	0.50	0.010	1.64	***	0.60	0.010	1.81	***
Capability								
geolocation	0.01	0.014	1.01		0.07	0.009	1.07	***
microphone	0.28	0.277	1.32	***	0.30	0.011	1.34	***
camera	0.32	0.315	1.37	***	0.30	0.011	1.35	***
User Action								
denied	-0.18	0.011	0.84	***	-0.30	0.010	0.74	***
dismissed	-0.19	0.010	0.83	***	-0.25	0.010	0.78	***
ignored	-0.20	0.010	0.81	***	-0.26	0.010	0.77	***
Had prior user interaction	0.04	0.008	1.04	***	0.05	0.008	1.06	***

5.5 RQ5: Who Benefits

Finally, as the last question of questionnaire 2, we asked participants who they thought would benefit from access to the requested capability. Based on existing work, we hypothesized that if users felt that they would benefit from capability access, they would be more likely to allow access. We find that perceiving a benefit for oneself is indeed more prevalent when making accept decisions (see Table 12 and Figure 5, 23-26% less likely to indicate self-benefit when not accepting). Additionally, respondents indicate perceiving more self-benefit when websites requested camera and microphone (34% and 35% more likely to indicate self-benefit); this suggests that, when these capabilities get requested, respondents can imagine fewer ways for using a camera or microphone other than for their benefit. Again, having had a prior user interaction only shows a very small effect on perceiving a benefit for oneself.

5.6 RQ6: The Role of User Interaction

For this RQ, we want to understand the impact of prior user interaction on the permission prompt outcomes we measured in this study. First of all, Section 3 shows that the 14.8% of prompts preceded by a user interaction had higher allow rates in general (29.8% with vs. 11.7% without). This effect is most pronounced for the geolocation capability (27.4% with vs. 9.0% without).

Table 13: Results of a logistic regression using whether or not respondents cited at least one rational reason for their decision as the dependent variable.

	At least one rational reason			
	Log odds	Std. Error	Odds	sig.
(Intercept)	0.72	0.010	2.06	***
Capability				
geolocation	0.02	0.009	1.02	*
microphone	-0.09	0.011	0.91	***
camera	-0.08	0.011	0.93	***
User Action				
denied	0.07	0.010	1.07	***
dismissed	-0.06	0.010	0.94	***
ignored	-0.21	0.010	0.81	***
Had prior user interaction	-0.01	0.008	0.99	
Is sure	0.07	0.008	1.07	***

The logistic regressions on annoyance, ease of decision-making and citing at least one rational reason shown in Tables 6, 7, and 13 also include having a prior user interaction as one independent variable. They show that prior interaction only significantly contributes to not feeling annoyed: participants who had a prior user interaction were 15% more likely to not feel annoyed. For the other outcomes, the presence of a user interaction is not a significant factor. Additionally, the two subsections above already show that

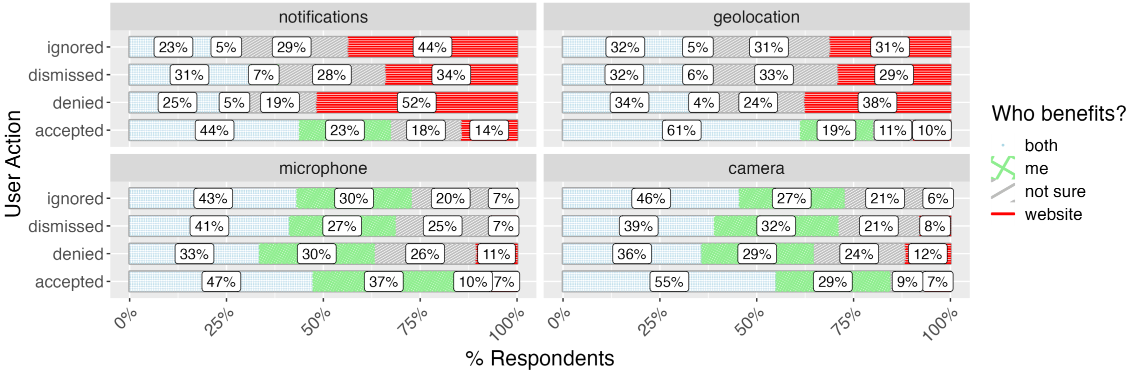

Figure 5: Respondents' answers to the question who would benefit from allowing access to a capability (Q3 in questionnaire 2) across the requested capabilities and action they took on the permission prompt.

there is only a very small effect of a prior user interaction on feeling sure and perceiving a self-benefit.

6 LIMITATIONS

Our work is limited in several ways. First of all, we were only able to investigate web permission usage and behavior on Chrome. While Chrome is the most popular desktop browser at the time of writing [24], there may be reasons why particular types of users turn to other browsers and could thus behave systematically different. Additionally, as discussed in Section 2, prompt UIs differ between browsers. Our results thus cannot be directly generalized to other browsers, especially those that are not based on Chromium.

Second, we focused this first investigation on user behavior and sentiment on Chrome for desktop platforms (Windows, macOS, Linux, and ChromeOS). Prior work has shown that at least the type of websites users visit are different on mobile devices [22]. Expanding this investigation to mobile versions of Chrome is planned as future work.

Third, we only collected telemetry and showed experience sampling prompts to users who have not opted-out of the "Help improve Chrome's features and performance" setting. Arguably, users who disable this setting may have different privacy attitudes and may thus react differently to permission prompts. Furthermore, telemetry collection and experience sampling were triggered when a website asks for permissions. Thus, websites visited more frequently by users who have not previously made a permanent permission decision were also more likely to contribute more data points to our sample. Investigating to what extent certain types of websites impact permission sentiment and behaviors is subject of future work.

Fourth, our approach was subject to self-selection bias. Users already annoyed by permission prompts may have been less likely to respond to our experience sampling prompt. In the light of the very low response rates we typically receive with such questionnaires in Chrome, this is a considerable concern. Due to the short and privacy-preserving nature of our experience sampling questionnaire, we unfortunately also do not know how representative the sample we obtained is when it comes to demographic properties such as age or gender. Additionally, we were only able to offer a limited set of reasons to users in the experience sampling questionnaires. It is possible that we are missing additional nuance captured by the additional items used by Bonné et al. [6] and Cao et al. [7]. However, despite these substantial limitations and thus a possibly imperfect representation of the user population or reduced nuance in captured reasons, we believe that collecting data in situ was worthwhile for a security UI that is very brief and contextual.

Finally, we focused on the most popular capabilities currently available on the web. As these also have been around for more than a decade, it stands to reason that browser makers, web developers, and users are familiar with them. The same is not necessarily true for other, more recently introduced, permission-gated capabilities. We thus cannot generalize beyond the four permission types we looked at in this study.

7 DISCUSSION

In the following subsections, we discuss what we believe to be the core findings of this work. As there may be substantial self-selection bias in our sample, we focus on findings between capabilities and user actions.

7.1 Permission Prompts are More Annoying When Users Do Not Allow

We find that the level of annoyance varies between requested capabilities and the decisions respondents took. They reported feeling more annoyed when they did not allow the requested access, especially for notification and geolocation capabilities. This implies that permission prompts are particularly annoying when one doesn't want the website to have that capability and the website is thus asking at an inopportune moment or without good reason. This is corroborated by allow actions being twice as common after a user interaction on the page, as well as being sure why the website is asking and a perceived self-benefit being associated with allowing access.

We believe this means that it needs to be easier for website developers to find better moments to ask for permissions, for example, by relying on users to trigger the use of a capability. This would also ensure that capability accesses unwanted by users are never prompted for. Additionally, our findings imply that developers should be encouraged to provide more contextual information before asking for permission. Both of these recommendations are mirrored in proposals for the Web (e.g. [21, 29]) and guidelines on other platforms, for example, on Android [13]. However, as noted at the beginning of the paper, the web does not rely on review processes and app stores to enforce such guidelines, and web developers thus need to be nudged towards better practices. Making such best practices at least easier to find and ideally designed into the capability APIs should be subject of future work.

7.2 Decision-Making on Permission Prompts is Problematic in Some Cases

While a majority of respondents cite at least one rational reason for doing what they did, two common problematic reasons surfaced: 9.1% of respondents who allowed capability access stated that they believed the website wouldn't work or wouldn't let them in altogether if they didn't allow, which is indicative of a feeling of being forced towards allowing access. This could be explained by a range of website behaviors, from simply using language that gets perceived as too strong to outright abusive behavior, where the website gates content on getting access to a powerful feature it wants to use for nefarious purposes, such as sending spammy notifications. Additionally, 30.4% of respondents stated that they did not notice the prompt and thus ended up ignoring it, which we discuss further in Section 7.5 below.

Decision-making also appears to be different between capabilities: reasons for making decisions on camera or microphone permission were in several cases selected at different frequencies for notification and geolocation access. Respondents also indicated feeling more sure why a website was asking for camera or microphone, and more respondents perceived a self-benefit when these capabilities are used. Future work should ensure that these capabilities are investigated separately, and system designers can also consider using different approaches for different capabilities.

7.3 Behavior and Decision-Making Appears Different from Mobile Platforms

We find that prior work on Android [6, 7] reported higher deny rates than we find for desktop Chrome. This is not surprising, given that Chrome's permission prompt offers ignore and dismiss as additional non-allow actions (see below). However, we also find substantially lower allow rates, especially for the geolocation (86% granted on Android vs. 27% and 9% granted with and without prior user interaction on desktop Chrome respectively) and camera (84-86% granted on Android vs. 70% and 67% granted with and without prior user interaction on desktop Chrome respectively) capabilities.

In addition, comparing the frequencies with which certain reasons were cited for decision-making, we find that functionality-related reasons were cited less frequently for allowing on Chrome desktop (47%) than on Android (68%). This supports the hypothesis that contextual information is less available on the web than on Android. Similarly, 23% stated that they wanted the permission prompt to go away when denying, while only 13% selected this reason in Bonné et al.'s Android study [6]. This was mostly driven by geolocation and notification prompts, supporting that those more frequent prompts cause more challenges for users than other types.

In sum, we hope our findings spark an interest in comparing permission systems across platforms, as our data suggests there are substantial differences in user behavior and decision-making when seemingly minor aspects of the prompt UI are modified. A particularly interesting piece of future work could be to run a study comparing the efficacy of permission systems between platforms.

7.4 Availability of Contextual Information and a Perceived Self-Benefit is Associated with Allowing

Supporting findings from prior work, we find that respondents who allowed capability access were much more likely to feel sure why the website was asking for access and also that there was a benefit for themselves. The delta in these perceptions between allow and non-allow actions suggests that the absence of such information is associated with respondents not wanting to grant the access. Additionally, the significant difference in sureness and perceived self-benefit, between notifications and geolocation on the one hand and camera and microphone on the other, suggests that contextual information is less available for notification and geolocation use cases on the web.

Given that most permission prompts on the web are not accepted, we interpret these findings to mean that website developers need substantial help with making sure they are requesting permission with sufficient contextual information and in situations where these capabilities bring value to their users. Identifying when it is helpful or detrimental to ask for access to the various capabilities on the web should be subject of future work.

7.5 Being Able to Ignore and Dismiss Permission Prompts is Useful

Ignoring and dismissing are the most common actions for the four most common permissions on the web. Contrary to a naive intuition, our results suggest that this is not primarily because users want to get the prompts out of the way. We find that 67% and 51% of our respondents cite rational reasons for dismissing and ignoring a prompt, respectively. Besides functionality-related reasons, 22% and 17% of respondents stated that they wanted to decide later, waiting to see if the capability is really necessary. The prominent option to dismiss via the "x" button in the top right corner of the prompt, as well as the option to easily ignore the prompt as it is neither modal nor visually very salient, appears to provide substantial value to Chrome's users.

We believe facilitating such actions should be considered for other permission UIs as well. However, we do find that 30.4% of respondents indicate that they ignore permission prompts because they did not notice the prompt. It stands to reason that at least some of these respondents may have later experienced functionality issues with the website. Understanding these situations better could be subject of future work.

In sum, there appears to be a delicate balance when offering an ignore option, that could also be improved in Chrome. The Chrome team is working on a draft proposal to the W3C [21] about ways to make permission prompts more visually salient without increasing annoyance and while maintaining the value of being able to ignore a website asking for a permission.

As a side note, we showed that temporary *deny* options (i.e. ignoring and dismissing) are popular on permission prompts in Chrome. However, there currently is no temporary *allow* option available to users. It seems reasonable to think that some users may want to temporarily allow a capability as well, for example to see if there is value in allowing long-term access. To that end, Chrome is in the process of adding an "allow this time" option, which began to roll out in August 2023 [9].

7.6 Prior User Interaction is Associated with Allowing and Being Less Annoying

Our findings suggest that prior user interaction with a website has some effects on respondents' behavior and perceived annoyance, but not on ease of decision-making or citing rational reasons for decisions. Telemetry showed that Chrome users are substantially more likely to allow permission requests with prior user interaction than without: 20% vs. 10% and 27% vs. 9% for notification and geolocation, respectively. Respondents to the experience sampling questionnaire were 15% more likely to not feel annoyed by the prompt after they had an interaction with the page. However, there was no effect at all for feeling that it was easy to make a decision or for citing a rational reason, as well as only a very small effect on feeling sure why the site is asking or benefiting from the capability.

The motivation behind the prior user interaction requirement, based on Mozilla's stated goals when introducing this heuristic for the notification permission, is to influence developer behavior in two ways: first, to incentivize developers to provide controls in the content area that allow users to initiate permission requests; second, to promote and facilitate providing additional contextual information for the request [19].

It is important to note that Firefox and Chrome, and thereby our analysis, use a definition of "prior user interaction" that corresponds to the "transient activation" state defined in the HTML standard [30], which is not necessarily associated with the permission request in a meaningful way. A permission prompt will be treated as having had prior user interaction if it is triggered within 5 seconds of a click or keyboard event on any part of the website's content; therefore, prior user interactions will include not only intentional interactions but also coincidental interactions that are unrelated to the permission request. Additionally, as the specification does not require any particular semantics for the piece of content the user interacted with, it seems likely that many developers chose the simplest way to satisfy the requirement and thus didn't provide much additional contextual information.

We speculate that the wide inclusion of interactions can lead to the higher allow rates observed after a user interaction as well as somewhat lower annoyance we found in our data. Even a simple button click would allow users to intentionally initiate a permission request. Conversely, the lack of developer-provided context could be why we don't detect a meaningful increase in ease of decision-making, citing rational reasons, feeling sure, and perceiving a self-benefit, as a button alone does not provide substantial contextual information. Based on this, we believe our data supports the hypothesis that, with its current implementation and at its current level of deployment to the web platform, the first design goal of the prior user interaction heuristic is achieved to some extent, while the second is not.

Dedicated future work could attempt to gate showing a permission prompt on more strongly associated user interactions, such as tying the triggering of prompts more closely to the ‘click‘ event handler of a button that provides additional contextual clues to users.

8 CONCLUSION AND NEXT STEPS

In this work, we describe user behavior and sentiment on web permission prompts in desktop Chrome. Given the differences in use cases as well as prompt UIs, we find that users appear to behave quite differently in comparison to the commonly studied mobile operating system permission prompts. A key difference is that deny rates are much lower on the web, given that other non-allow actions are available. We find some indications that users' reasoning about their decisions is also different, so we encourage dedicated, comparative research on permission systems between platforms.

For web permission prompts on desktop operating systems, we identify several areas for improvement. First, respondents rate permission prompts on the web as less annoying when prompts are allowed, implying that prompts where users do not find capability access necessary are more problematic. Thus, we believe the main challenge with permission prompts on the web lies in having websites make fewer unnecessary capability requests that users then choose to not allow, as such requests appear to be the main cause of annoyance. Put another way, we think web developers need to be encouraged to leverage best practices for requesting permissions, similar to those that are available for Android [13] and iOS [3].

Second, our data suggests that the ability to ignore and dismiss permission prompts is valuable for Chrome users, given that many respondents cite rational reasons for using these actions. There is, however, room for improvement, as 30.4% of respondents indicate they ignored the permission prompt because they did not notice it.

To address these two key issues, the Chrome team is currently exploring an alternate approach to permission prompts that may help nudge developers towards asking at more opportune times with more contextual information and alleviate the problem of users not noticing prompts [21].

We also find that several factors appear to moderate decision-making in line with findings in prior work. Having a user interaction with the website prior to the permission prompt makes respondents 15% more likely to not be annoyed and increases the rate of allow decisions up to three times (9% vs. 27% without and with prior interaction for geolocation). The likelihood of allow decisions is higher when respondents indicate being sure why the website is asking for the permission and when they perceive a benefit for themselves from allowing access to the desired capability. We also find that decision-making and sentiments were substantially different between the capabilities we investigated, with notifications and geolocation seeing higher non-allow rates and higher annoyance when not allowing. This suggests that future investigations of permission behaviors should look at different capabilities separately.

Finally, our research supports the recommendation that permissions should be asked for when there's a high chance of users actually granting them, as that makes prompts less annoying. We also find that the current prior user interaction heuristic via the transient activation mechanism may not be ideal to approximate a good moment to ask for permissions, due to limited developer incentives to adhere to it and to provide additional contextual information.

ACKNOWLEDGMENTS

We would like to thank Balazs Engedy for his in-depth advice and review of this paper and Florian Jacky for implementing the experience sampling infrastructure. We are also grateful for Mike West, Emanuel von Zezschwitz, Chris Thompson, Serena Chen, Meggyn Watkins, Adriana Porter Felt as well as the anonymous CHI reviewers and ACs helping us to improve the manuscript.

REFERENCES

[1] Chaitrali Amrutkar and Patrick Traynor. 2012. Short Paper: Rethinking Permissions for Mobile Web Apps: Barriers and the Road Ahead. In *Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*. ACM, New York, NY, USA, 15–20. https://doi.org/10.1145/2381934.2381939.

[2] Android Developers. 2023. Notification runtime permission. https://developer.android.com/develop/ui/views/notifications/notification-permission. Last accessed: 2023-07-25.

[3] Apple. 2023. App Store Review Policy – Privacy. https://developer.apple.com/app-store/review/guidelines/#privacy. Last accessed: 2023-07-13.

[4] Igor Bilogrevic, Balazs Engedy, Judson L. Porter III, Nina Taft, Kamila Hasanbega, Andrew Paseltiner, Hwi Kyoung Lee, Edward Jung, Meggyn Watkins, PJ McLachlan, and Jason James. 2021. "Shhh...be quiet!" Reducing the Unwanted Interruptions of Notification Permission Prompts on Chrome. In *USENIX Security*. USENIX Association, Santa Clara, CA, 769–784. https://www.usenix.org/conference/usenixsecurity21/presentation/bilogrevic.

[5] Rainer Böhme and Jens Grossklags. 2011. The Security Cost of Cheap User Interaction. In *New Security Paradigms Workshop (NSPW)*. ACM, New York, NY, USA, 67–82. https://doi.org/10.1145/2073276.2073284.

[6] Bram Bonné, Sai Teja Peddinti, Igor Bilogrevic, and Nina Taft. 2017. Exploring decision making with Android's runtime permission dialogs using in-context surveys. In *SOUPS*. USENIX Association, Santa Clara, CA, 195–210. https://www.usenix.org/conference/soups2017/technical-sessions/presentation/bonne.

[7] Weicheng Cao, Chunqiu Xia, Sai Teja Peddinti, David Lie, Nina Taft, and Lisa Austin. 2021. A Large Scale Study of Users Behaviors, Expectations and Engagement with Android Permissions. In *USENIX Security*. USENIX Association, Santa Clara, CA, 803–820. https://www.usenix.org/conference/usenixsecurity21/presentation/cao-weicheng.

[8] Chrome Developers Blog. 2018. Unlocking new capabilities for the web. https://developer.chrome.com/blog/capabilities/. Last accessed: 2023-07-05.

[9] Chrome Developers Blog. 2023. One-time permissions in Chrome. https://developer.chrome.com/blog/one-time-permissions/. Last accessed: 2023-11-20.

[10] Yusra Elbitar, Michael Schilling, Trung Tin Nguyen, Michael Backes, and Sven Bugiel. 2021. Explanation beats context: The effect of timing & rationales on users' runtime permission decisions. In *USENIX Security*. USENIX Association, Santa Clara, CA, 785–802. https://www.usenix.org/conference/usenixsecurity21/presentation/elbitar.

[11] Adrienne Porter Felt, Serge Egelman, and David Wagner. 2012. I've Got 99 Problems, but Vibration Ain't One: A Survey of Smartphone Users' Concerns. In *Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*. ACM, New York, NY, USA, 33–44. https://doi.org/10.1145/2381934.2381943.

[12] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android Permissions: User Attention, Comprehension, and Behavior. In *SOUPS*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/2335356.2335360.

[13] Google Play Console Help. 2023. Declare permissions for your app. https://support.google.com/googleplay/android-developer/answer/9214102. Last accessed: 2023-07-13.

[14] Marian Harbach, Igor Bilogrevic, Enrico Bacis, Serena Chen, Ravjit Uppal, Andy Paicu, Elias Klim, Meggyn Watkins, and Balazs Engedy. 2024. Don't Interrupt Me – A Large-Scale Study of On-Device Permission Prompt Quieting in Chrome. In *NDSS*. The Internet Society, San Diego, CA, 14 pages. https://dx.doi.org/10.14722/ndss.2024.24108.

[15] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions. In *CHI*. ACM, New York, NY, USA, 2647–2656. https://doi.org/10.1145/2556288.2556978.

[16] Mohammadreza Hazhirpasand, Mohammad Ghafari, and Oscar Nierstrasz. 2020. Tricking Johnny into Granting Web Permissions. In *International Conference on Evaluation and Assessment in Software Engineering (EASE)*. ACM, New York, NY, USA, 276–281. https://doi.org/10.1145/3383219.3383248.

[17] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. 2012. A Conundrum of Permissions: Installing Applications on an Android Smartphone. In *Financial Cryptography and Data Security*. Springer, Berlin, Heidelberg, 68–79. https://doi.org/10.1007/978-3-642-34638-5_6.

[18] Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L. Mazurek, and Jeffrey S. Foster. 2017. User Interactions and Permission Use on Android. In *CHI*. ACM, New York, NY, USA, 362–373. https://doi.org/10.1145/3025453.3025706.

[19] Mozilla. 2019. Restricting Notification Permission Prompts in Firefox. https://blog.mozilla.org/futurereleases/2019/11/04/restricting-notification-permission-prompts-in-firefox. Last accessed: 2023-07-13.

[20] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[21] Andy Paicu, Balazs Engedy, Marian Harbach, Serena Chen, Penelope McLachlan, Thomas Nguyen, and Kamila Hasanbega. 2023. Page Embedded Permission Control (PEPC) Explainer. https://github.com/WICG/PEPC/blob/main/explainer.md. Last accessed: 2023-11-17.

[22] Kimberly Ruth, Aurore Fass, Jonathan Azose, Mark Pearson, Emma Thomas, Caitlin Sadowski, and Zakir Durumeric. 2022. A World Wide View of Browsing the World Wide Web. In *Internet Measurement Conference (IMC)*. ACM, New York, NY, USA, 317–336. https://doi.org/10.1145/3517745.3561418.

[23] Donald Sharpe. 2015. Chi-square test is statistically significant: Now what? *Practical Assessment, Research, and Evaluation* 20, 1 (2015), 8. https://scholarworks.umass.edu/pare/vol20/iss1/8/.

[24] Statcounter. 2023. Desktop Browser Market Share Worldwide. https://gs.statcounter.com/browser-market-share/desktop/worldwide. Last accessed: 2023-07-13.

[25] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *CHI*. ACM, New York, NY, USA, 24 pages. https://doi.org/10.1145/3544548.3581060.

[26] Christopher Thompson, Maritza Johnson, Serge Egelman, David Wagner, and Jennifer King. 2013. When It's Better to Ask Forgiveness than Get Permission: Attribution Mechanisms for Smartphone Resources. In *Symposium on Usable Privacy and Security (SOUPS)*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/2501604.2501605.

[27] Daniel Votipka, Seth M. Rabin, Kristopher Micinski, Thomas Gilray, Michelle M. Mazurek, and Jeffrey S. Foster. 2018. User Comfort with Android Background Resource Accesses in Different Contexts. In *Symposium on Usable Privacy and Security (SOUPS)*. USENIX, Baltimore, MD, 235–250. https://www.usenix.org/conference/soups2018/presentation/votipka.

[28] W3C. 2021. Secure Contexts. https://www.w3.org/TR/secure-contexts/. Last accessed: 2023-07-05.

[29] Mike West. 2023. Purposeful Permissions. https://github.com/mikewest/purposeful-permissions. Last accessed: 2023-11-17.

[30] WhatWG. 2023. HTML Living Spec – Tracking User Activation. https://html.spec.whatwg.org/multipage/interaction.html#tracking-user-activation. Last accessed: 2023-07-10.

[31] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android Permissions Remystified: A Field Study on Contextual Integrity. In *USENIX Security*. USENIX Association, Washington, D.C., 499–514. https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera.

[32] Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman. 2018. Contextualizing Privacy Decisions for Better Prediction (and Protection). In *CHI*. ACM, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173842.

A QUESTIONNAIRES

A.1 Questionnaire 1

Question text variables:

- *$capability = {"notifications", "location", "camera", "microphone"}*

Questions:

Q0. We are working on making it easier to control which websites can access [your] $capability. Help us by taking a short survey!

Q1. How annoying did you find having to make a decision on $capability access for this website?

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

Q2. To what extent did you feel interrupted by having to make a decision on $capability access for this website?

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

Q3. How easy or difficult did you find making a decision on $capability access for this website?

- Very difficult
- Somewhat difficult
- Neither difficult nor easy
- Somewhat easy
- Very easy

Q4. Thank you for helping to improve Chrome!

A.2 Questionnaire 2

Question text variables:

- *$capability = {"send(ing) you notifications", "your location", "your camera", "your microphone"}*
- *$action = {"allow(ed)", "dismiss(ed)", "ignore(d)", "[didn't|not] allow"}*

Questions (bold items highlight differences between items provided in Q1 between different conditions):

Q0. A website just asked for access to $capability. Help us improve how websites ask for access by taking this 1-minute survey.

Q1. When the website asked for access to $capability, why did you do what you did? Select all that apply. [randomized option order]

- *if user action is "dismiss"*
  - The website shouldn't need $capability
  - I want to be asked again later
  - I don't want to use any feature associated with $capability
  - I don't trust the developer enough to provide this information
  - I wanted the popup to go away
  - I don't remember what I did
  - Other (please specify): [not randomized]
- *if user action is "ignore"*
  - The website shouldn't need $capability
  - I want to be asked again later
  - I don't want to use any feature associated with $capability
  - I don't trust the developer enough to provide this information
  - I wanted the popup to go away
  - **I did not notice the website asking for/to $capability**
  - I don't remember what I did
  - Other (please specify): [not randomized]
- *if user action is "deny"*
  - The website shouldn't need $capability
  - **I can always allow it later if I change my mind**
  - I don't want to use any feature associated with $capability
  - I don't trust the developer enough to provide this information
  - I wanted the popup to go away
  - I don't remember what I did
  - Other (please specify): [not randomized]
- *if user action is "allow"*
  - **I want to use a feature that requires $capability**
  - **I trust the website's developer**
  - **I think the website won't work/let me in otherwise**
  - **Nothing bad will happen**
  - I wanted the popup to go away
  - **I won't be able to allow $capability later**
  - I don't remember what I did
  - Other (please specify): [not randomized]

Q2. How sure are you about why this website was asking for access to $capability?

- Not at all sure
- Slightly sure
- Moderately sure
- Very sure
- Extremely sure

Q3. Why do you think this website primarily asked for access to $capability? [randomized option order]

- For their own benefit
- For my benefit
- For both, their and my benefit
- I'm not sure

Q4. Thank you for helping to improve Chrome!

B ADDITIONAL FIGURES

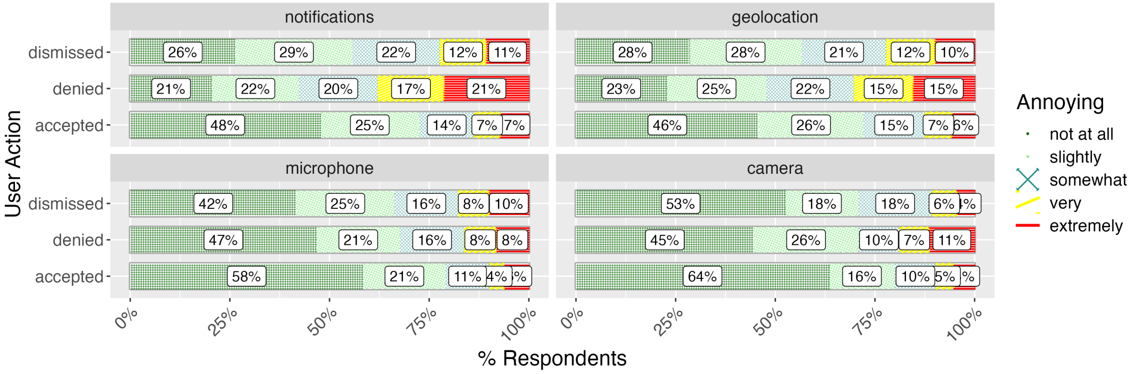

Figure 6: Respondents' answers to the question how annoying they found the permission prompt (Q1 in questionnaire 1) across the requested capabilities and action they took on the permission prompt.

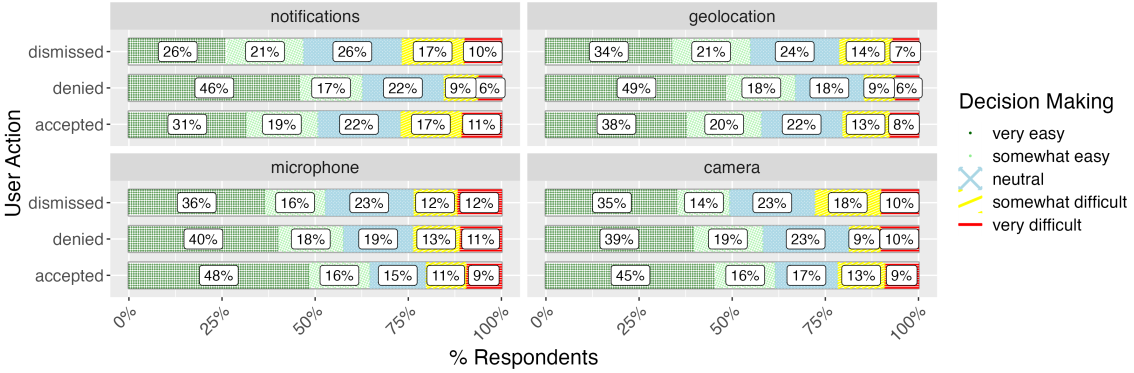

Figure 7: Respondents' answers to the question how easy or difficult it was to make a decision on the permission prompt (Q3 in questionnaire 1) across the requested capabilities and action they took on the permission prompt.