The Power of Words: A Comprehensive Analysis of Rationales and Their Effects on Users' Permission Decisions

Yusra Elbitar*†, Alexander Hart* and Sven Bugiel*
*CISPA Helmholtz Center for Information Security, †Saarland University
{yusra.elbitar, hart, bugiel}@cispa.de

***Abstract*—Rationales offer a method for app developers to convey their permission needs to users. While guidelines and recommendations exist on how to request permissions, developers have the creative freedom to design and phrase these rationales. In this work, we explore the characteristics of real-world rationales and how their building blocks affect users' permission decisions and their evaluation of those decisions. Through an analysis of 720 sentences and 428 screenshots of rationales from the top apps of Google Play, we identify the various phrasing and design elements of rationales. Subsequently, in a user study involving 960 participants, we explore how different combinations of phrasings impact users' permission decision-making process. By aligning our insights with established recommendations, we offer actionable guidelines for developers, aiming to make rationales a usable security instrument for users.**

I. Introduction

Imagine you open an app, and one of the following messages pops up: *"We would like access to your camera for the app to work correctly,"* or *"Please allow access to your camera. Without this permission, the app cannot scan your documents. We do not collect or transfer any personal data outside your phone."* Which permission request would you rather approve? Which would leave you feeling better informed, more satisfied, and more in control of your decision? In a world where app developers have the freedom but also the responsibility to clarify their permission requests to users, it is crucial that the wording and presentation of these explanations align with users' expectations.

In this paper, we explore *"rationales"*—the explanations behind runtime permission requests—by analyzing their designs, wording, and how different phrasings affect users' permission decisions. **Our first goal is to understand the current state of rationales in Android apps** (see Figure 1a). Despite numerous guidelines for developers on effective permission communication [5], [34], [53], the practical implementation and interpretation of these recommendations in the design and phrasing of real-world rationales remains uncertain. By manually analyzing 720 sentences and 428 screenshots of rationales collected from the top apps on Google Play, we unveil how app developers convey their permission needs in practice. This investigation establishes a baseline for evaluating future innovations and changes in this domain, potentially guiding the improvement of existing best practices for the benefit of both app developers and end-users.

Our exploration revealed the diversity of rationales with various patterns. We found that app developers phrase their rationales from different perspectives. Rationales can comprise one or multiple phrases. They can be specific or vague, positive or negative, and may include additional information as optional clauses. Regarding design, we observed diverse layouts with unique combinations of buttons, icons, and titles. Additionally, we observed patterns that occurred more frequently among developers, such as using a dialog to present a concise rationale, typically composed of a single phrase.

Finding a variety of different ways how app developers can convey permission needs to users naturally leads to **the second question of whether different rationale phrasings impact users' decision-making process** (see Figure 1b). Previous research in other fields has shown that linguistic variations can influence a variety of user decisions [1], [16], [29], [44], [60], [66]. Therefore, we break down rationales into their fundamental building blocks and examine how different combinations of these building blocks affect users' permission decisions and their perception of those decisions.

In an online user study with 960 participants, we gain insights into how users perceive and respond to rationales. We demonstrate that rationale phrasings alone significantly impact user choices regarding permissions and their assessment of these choices. When comparing the two rationales at the outset of this introduction, our study reveals that the second phrasing leads to a higher likelihood of granting permission, provides users with more informed decision-making, increases user satisfaction, and enhances the perception of being in control.

We observed that the natural variability in rationales is inherent. However, specific phrasings within rationales are essential for app developers to prioritize to improve the user's decision-making process. Finally, we compare our findings with available guidelines to create actionable recommendations for app developers, aiming to make rationales a usable security instrument for users.

Network and Distributed System Security (NDSS) Symposium 2025
24-28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-8-3
https://dx.doi.org/10.14722/ndss.2025.230544
www.ndss-symposium.org

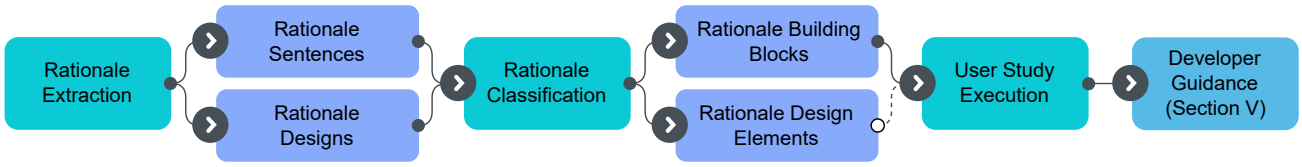

(a) Section III: Investigating Rationale Differences (b) Section IV: User Study

Fig. 1: Methodology of our exploration of rationales and their effect on users' permission decisions.

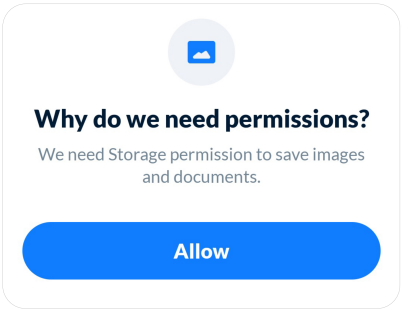

Fig. 2: Sample rationale from an Android app.

II. Related Work

The persistent challenge of aligning app developers' needs for their apps to function with users' willingness to grant permissions has been extensively studied. Previous research has revealed a disparity between the permissions that apps request and the users' expectations [8], [9], [11], [22], [41], [70]. This disparity creates a sense of discomfort and uncertainty among users, impacting their confidence in making a permission decision. Hence, there is a need for proper communication between developers and users.

App developers can use various instruments to communicate permission usage to users, including app descriptions, privacy policies, and, more recently, iOS Privacy Nutrition Labels [6] and Google Play's Data Safety Section [35]. Given the wealth of information in app descriptions, research has proposed techniques to automatically extract permission usage from text or code [51], [54], [55], [67], [73], which can then be used to measure the fidelity of the description in conveying permission usage to users [23], [25], [57]. The primary goal of iOS Privacy Nutrition Labels and Google Play's Data Safety Section is to streamline the content of lengthy and complex privacy policies. However, these instruments have their own limitations and challenges in terms of being correctly interpreted by both developers and users [15], [42], [46], [47], [58], [59], [74], [75]. As a result, a dedicated research area is focused on enhancing the usability of these privacy instruments [15], [48], [74], [75].

Another, more direct approach to explaining the need for permissions to users is through rationales. A rationale is a message associated with a runtime permission request. For instance, in iOS, the message *"We need Storage permission to save images and documents."* is directly integrated into the permission request, while in Android, it can be presented in a separate view, which developers can design as they like (see Figure 2). Although general guidelines exist for crafting rationales [5], [34], [53], prior research indicates that app developers often underutilize them, and when used, they often lack meaningful content [50], [65]. Given that app stores do not impose any specific regulations on how rationales should be presented to users, two crucial questions remain unanswered: (1) *how do app developers currently phrase and design rationales*, and (2) *how do different rationale phrasings influence users' permission decisions?*

Limited research has explored the influence of rationales on user decision-making. Previous studies have shown that users benefit from receiving additional information in permission requests [61], [69], [72]. Specifically, earlier research has indicated that including rationales increases the likelihood of users granting permissions [65] and positively affects how users assess their permission choices [19]. Furthermore, similar findings exist in other domains. For instance, in the field of two-factor authentication (2FA), personalized messages have been found to enhance user adoption of 2FA [29]. Additionally, the terminology used in app descriptions significantly shapes user perceptions of security within secure messaging apps [1], [16]. However, it is crucial to emphasize that the specific effects of rationales have yet to be investigated.

While the influence of rationale phrasing on users' permission decision-making has received limited attention, the significance of language in decision contexts is well-established in various fields. One such phenomenon is the framing effect [66], a concept extensively studied in psychology [63]. The framing effect reveals how decisions and judgments can be swayed by how information is presented or framed [44], [66]. Other studies have shown that subtle wording differences can significantly influence customers' attitudes toward brands [60]. By drawing connections between these diverse findings, it is reasonable to conclude that users deciding on permissions affecting their privacy may also be influenced by variations in the rationale prompt, whether or not these variations provide new information or are merely linguistic differences.

III. Investigating Rationale Differences

While numerous recommendations exist for developers about how to communicate permission needs to users [5], [34], [53], the actual interpretation and implementation of these recommendations in rationales remain unclear. The following exploratory analysis aims to unveil how permission

needs are communicated in practice. To this end, we first crawled the most relevant apps from Google Play. Next, we extracted rationales and manually labeled them in a bottom-up coding process. Finally, we extracted dimensions along which developers' implemented rationales typically differ.

Between 08/2021 and 04/2022, we continuously crawled the top 50 apps in every category on Google Play. This effort produced a dataset containing 11,500 unique APKs, considering solely the most recent app versions. For our exploratory analysis, we narrowed this dataset to apps that requested at least one runtime, i.e., dangerous protection level, permission in their manifest file. This yielded a final selection of 9,489 APKs, set for in-depth exploration.

A. Rationale Extraction and Classification

Our next step was to extract rationales. Rationales can come in various forms and may not be immediately recognizable within the app's UI elements. However, we saw this as a valuable starting point for locating rationales because Google recommends that developers specify all text elements in the strings.xml files of Android apps [32]. Given the challenge of identifying rationales due to their scarcity, we developed an initial classifier using SpaCy [21]. This classifier was trained on a labeled dataset from our prior work [19], which included 450 rationales and 250 non-rationales. We employed this classifier to filter out non-rationales from all sentences extracted from the string.xml files, resulting in a dataset of 55,000 unique sentences that had the potential to be rationales.

During the initial classifier assessment, we manually checked 3,945 of these sentences. Once labeled, cleaned, and cleared of duplicates, these sentences comprised 801 rationales and 892 non-rationales set aside for threshold optimization and evaluation. The remaining sentences underwent annotation using Prodigy [20], a tool that employs active learning, selectively prompting users to annotate sentences that the classifier struggles with. In total, 1,500 sentences were annotated—777 as rationales and 723 as non-rationales—forming our balanced training dataset. Our final classifier achieved a precision of 0.99, a recall of 0.84, and an F-score of 0.91.

Applying this classifier to the strings.xml files, we found 35,737 unique rationale sentences across 6,524 apps. Our investigation then broadened to analyze the overall context in which these rationales appeared through dynamic analysis. This involved exploring design aspects and phrases that, while not rationales on their own, carry significance, such as *"See how we protect your privacy, tap here."* or *"We will not collect your personal information."*

To execute dynamic analysis on our dataset, we employed an approach inspired by prior research [12], [13]. Our analysis was conducted simultaneously on four Android emulators (API level 30), with a timeout of 12 minutes per app. We started with decoding the APK and rendering its activities accessible to external entities by setting *"exported=true"* for each activity within the manifest file. We used Apktool [39] to decode and repackage APKs. Then, we launched each activity through adb shell commands and automatically navigated through its interactive components. Permission requests encountered during the launch were intentionally denied and followed by an activity restart. This was essential to activate rationales that emerge specifically after at least one permission denial [30]. Subsequently, we dumped the XML layout of the current activity and used our classifier to analyze all sentences within it for potential rationales. When detected, we captured a screenshot of the respective activity.

Given the complexity of Android activity layouts, potentially including hidden elements, we harnessed UIAutomator [36] to interact with interactive components. Our approach extended to testing swipe gestures, especially relevant when dealing with onboarding processes devoid of clickable *"next"* buttons. Interactions that led to layout changes within the same activity prompted a subsequent rationale check.

We successfully explored 3,818 of the 6,524 apps, encompassing partially analyzed apps due to timeouts (346 APKs). Unfortunately, analysis was unfeasible for 1,486 apps, as their activity initiation necessitated extra parameters ascertainable only through time-consuming static analysis [12]. Apktool could not repack the remaining 1,220 apps.

In total, we collected 2,953 screenshots. We filtered out screenshots depicting cookie notifications, update messages, privacy policies, and other instances erroneously included by our classifier due to resemblances with rationales. While some of these instances did contain rationales, they were often embedded within privacy policies, which were beyond the scope of this study. We also removed duplicate rationales within apps. This curation process resulted in a total count of 1,054 unique rationale screenshots from 709 distinct apps.

We analyzed both the content and design of rationales using inductive and axial coding until saturation. Initially, our dataset consisted of 35,737 unique text-based rationales and 1,054 rationale screenshots. Given the substantial number of rationales to code manually, our methodology involved assigning categorical labels (aka. codes) to rationale messages and designs. In this process, two independent researchers created codes for the same batch of rationales, which were then collaboratively discussed until an amalgamated codebook was formed. The final codebook was subsequently used to code 720 randomly chosen text rationales and 428 screenshots, creating our sample dataset. The entire coding process took two independent researchers four weeks to complete. The final codebook can be found on the Open Science Framework [18].

Next, we will outline our findings on the diversity of rationale phrasing and design. We will present the occurrence percentages of each label in two datasets: our manually labeled rationale sentences ($N = 720$, marked as $\%^{m}$) and our manually labeled screenshots ($N = 428$, marked as $\%^{s}$).

1) Rationale Building Blocks: This section focuses on the structure and phrasing of rationale sentences.

Functionality. In our sample, rationales varied between providing specific information about permission use and remaining vague. At times, they failed to reveal the functionality requiring permission. In our manual analysis, when a rationale

pertains to a particular feature within the app (e.g., searching for gas stations), we considered it specific (58%^{m} 69%^{s}). Conversely, a functionality that merely indicates that permission is necessary would be labeled unspecific (12%^{m} 14%^{s}). In some cases, no functionality was indicated (30%^{m} 17%^{s}).

Articulation. There are two ways to convey permission-enabled functionality to users. One approach is positive phrasing (79%^{m} 83%^{s}), where granting permission enables the functionality. The other is negative phrasing (21%^{m} 17%^{s}), where lack of permission results in the functionality being unavailable. For instance, positive phrasing would be something like *"Camera permission is needed to use the scanner,"* while its negative counterpart would be *"Unable to use the scanner without the camera permission."*

Permission Type. We observed that rationales typically specify the permission they pertain to (83%^{m}). Within our dataset, the most prevalent permissions were location (32%^{m}), storage (29%^{m}), camera (19%^{m}), and microphone (8%^{m}). Additionally, rationales might encompass multiple permissions (9%^{m} 7%^{s}), falling into two categories: either several permissions collectively enable the same functionality, as illustrated in Figure 3a, or each permission individually facilitates different functionalities, as depicted in Figure 3b.

Perspective. A rationale can be phrased from one or multiple perspectives. One perspective addresses the user and prompts them to take action. Alternatively, rationales can emphasize the app's necessity for permission to execute specific functions. Another approach involves highlighting the permission's role in facilitating particular functionalities.

- **User Perspective.** When a rationale is framed from the users' perspective, it prompts them to take steps toward granting permission. The nature of these actions can vary in terms of specificity. In its simplest form, it might ask the user to *"please grant permission."* Alternatively, it could guide where to click to grant permission, like *"click allow to grant permission"* or *"grant permission in the next screen."* In our dataset, most rationales contained a phrase from the users' perspective (67%^{m} 59%^{s}).
  When permission is blocked, users must take extra steps to grant it. Rationales associated with this situation usually direct users to device settings, e.g., *"please grant permission from settings"* (24%^{m} 18%^{s}). For more precise instructions, these directions are frequently laid out in a step-by-step manner, utilizing commas (,,,), operators that are used as arrows (>>>), or numbered lists (1. 2. 3.). For example, a direction might read: *"Tap settings > go to app info > permissions, then allow permission."*
  We have identified two distinct approaches to how a rationale addresses the user. The first involves employing imperative commands with words like *"grant," "turn on," "allow,"* and *"enable."* This kind of rationale can come across as demanding. However, adding the word *"please"* to the request introduces a more polite tone. We found that around half of the rationales with the user perspective were politely phrased (43%^{m} 40%^{s}). The second approach directly addresses the user using phrases like *"you must," "you will need to,"* and *"you have to."* Furthermore, the users' perspective can also be conveyed in negative sentences like *"you denied permission."*
- **App Perspective.** Another common perspective in rationales is that of the app itself. (30%^{m} 37%^{s}). In this phrasing style, we encountered several variations. The app explicitly expresses its need for permission, often stating *"this app needs permission"* or *"{app_name} needs permission."* Sometimes, specific functionalities of the app need permission, as seen in *"scanning QR codes needs permission."* A more polite approach would be, *"this app would like permission,"*. However, we did not come across this polite variation very often in rationales that included a phrase from the app's perspective (2%^{m} 3%^{s}). The app's perspective is also reflected in negative sentences such as *"the app does not have access to permission."*
- **Objective Perspective.** When the need for permission is communicated passively, such as in the phrases *"permission is needed"* or *"permission is used,"* the rationale takes on an objective tone (27%^{m} 30%^{s}). We also found instances where the request to grant permission is presented in a passive form, as evident in *"permission must be granted."* Additionally, when expressed negatively, examples include *"permission denied"* or *"without this permission, the app cannot function."*

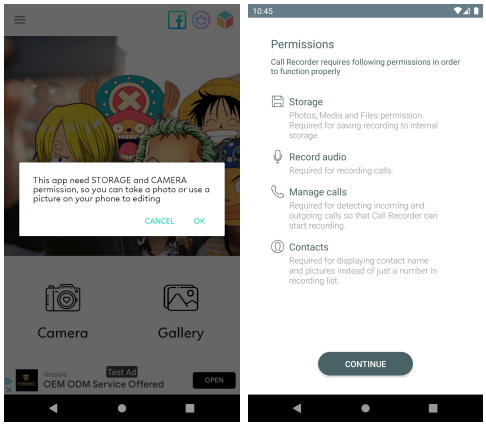

(a) (b)

Fig. 3: (a) Multiple permissions enable the same functionality. (b) Each permission enables a different functionality.

User-Centric. A rationale can center around the user. This can be achieved by either aligning the permission with the user, as seen in *"access to your location"* (29%^{m}) or by emphasizing the benefits of granting permission for the user, like *"to search for gas stations near you"* (18%^{m} 17%^{s}).

Optional Building Blocks. Apart from the above core component of a rationale, we discovered that a rationale can encompass one or several optional building blocks:

- **Empower with Control (control):** One such element involves empowering users to manage their permission choices at any point, such as *"you can change permissions from device settings anytime"* (1%^{m} 3%^{s}).

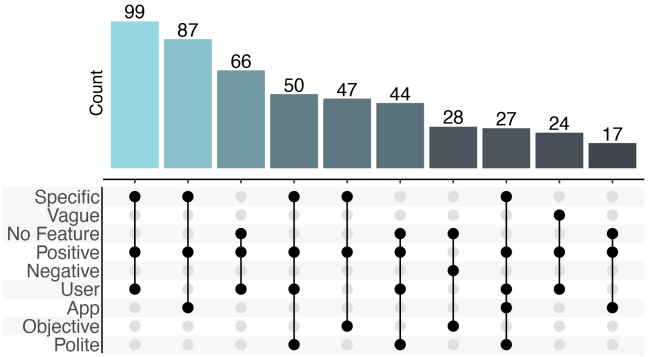

Fig. 4: Most common rationale building block combinations (dataset: sentences $N = 720$).

- **Mitigate (Mis)use (guarantee):** The rationale can reassure users that permission is used for a specific purpose, such as *"we will only use your permission for smart tracking,"* or clarify it will not be used for other reasons, like *"we do not track your location"* (2%^{m} 6%^{s}).
- **Offer Alternatives:** An alternative solution can be integrated into the rationale, giving users an option if they choose not to grant permission. For example, a rationale can include the phrase *"alternatively, you can specify your location manually"* (2%^{m} 3%^{s}).
- **More Information (more):** A link to more information or the app's privacy policy can be included (1%^{m} 6%^{s}).
- **Prompt Engagement:** The rationale can incorporate a question, encouraging users to actively decide on granting permission, such as *"do you want to allow this permission?"* or to confirm denial, like *"are you sure you want to deny this permission?"* (3%^{m} 6%^{s}).

Rationale Timing. Rationales can be linked to first-time permission requests (type I). They can also belong to previously denied requests (type II), often mentioning that the app lacks permission. Additionally, some rationales are connected to permission requests that have been blocked (type III). This can happen if the user denies permission multiple times within the same app life-cycle or if they have selected the *"never ask again"* option for pre-Android 11 apps. These rationales may include phrases suggesting the user can grant permission from the device settings.

Fig. 5: Labeled rationale with two phrases.

Multiple Phrases. While rationales are typically short, in some cases, they consist of two or three phrases (34%^{m}, 33%^{s}). When rationales consist of multiple phrases, the most common combination involves a phrase from the user's and one from the app's perspective, as demonstrated in the labeled rationale in Figure 5 (cf. Figure 6). Additionally, rationales may include optional building blocks.

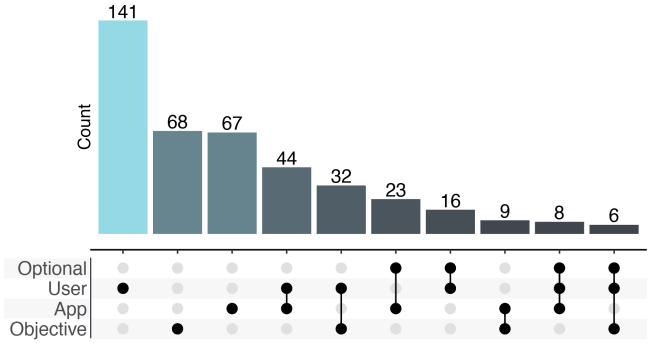

Fig. 6: Most common combinations of optional building blocks and perspectives (dataset: screenshots $N = 428$).

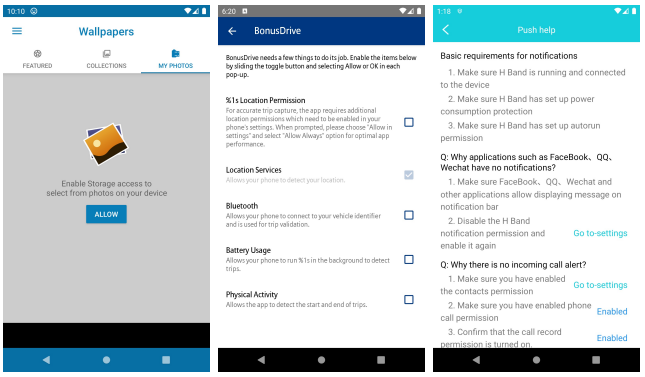

(a) Main (b) Settings (c) Help

Fig. 7: Different rationale embedding points in the app.

Summarizing our findings regarding rationale phrasing (illustrated in Figure 4), we observe that most consist of a single phrase. Our analysis reveals that perspectives predominantly stem from the user, trailed by those from the app and objective. Additionally, a substantial portion of these rationales specifies a functionality, highlighting that granting permission will enable this functionality, a strategy we term positive articulation.

2) Rationale Design Elements: Next, we will outline our findings concerning the design aspects surrounding a rationale.

Presentation. In exploring 428 rationale screenshots, we discovered that rationales can take on diverse formats, as shown in Figure 8. The most prevalent choice is a dialog-style rationale (59%^{s}), followed by fullscreens (23%^{s}), which were sometimes included in the onboarding process, embedded forms (10%^{s}), and banners (8%^{s}).

- **Dialog.** Dialogs typically follow standard Android styling [31], allowing developers to use them without modifications, as seen in Figure 8a. However, they can also be customized to resemble fullscreen rationales.
- **Fullscreen.** Fullscreens lack a default style, with their appearance based on the developer's preferences. Their spacious layout allows for more detailed information (Figure 8b). A variation includes integrating the rationale into the app's onboarding process (44%^{s}), typically shown at first launch unless skipped.

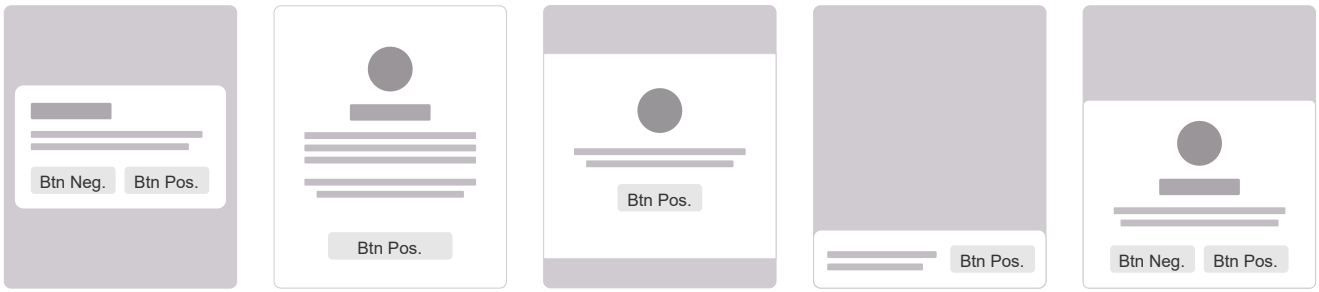

(a) Dialog (59%) (b) Fullscreen (23%) (c) Embedded (10%) (d) Banner Modeless (4%) (e) Banner Modal (4%)

Fig. 8: The different rationale presentations. Pos.=Positive, Neg.=Negative.

- **Embedded.** Embedded rationales are integrated into the app's main screens. They often replace permission-protected content until permission is granted, as shown in Figure 8c (cf. Figure 7a). Additionally, they can be part of the settings screen (see Figure 7b). This rationale stands out due to its interactive nature. It allows users to activate or deactivate permissions using buttons, checkboxes, or toggle switches, visually indicating the permission's status. Alternatively, a rationale can be integrated into a help or troubleshooting screen (see Figure 7c).
- **Banner.** There are two forms of Banners. The first form is narrow and includes the rationale message and a button, as in Figure 8d. This style is called modeless, meaning it does not interrupt the user's ongoing activity and usually follows the standard Android appearance of banners. Half of the banners followed this form (50%^{s}). The other half took up more space and were modal (50%^{s}), meaning they require user interaction, as in Figure 8e.

Buttons. A rationale can feature one or multiple buttons. In cases where only one button is available (54%^{s}), its actions vary. Usually, when this button is labeled with terms like *"ok," "allow," "grant," "enable," "next," "continue,"* or *"proceed,"* it serves a positive function and triggers the display of a permission request. However, this positive button may guide users to the settings screen for rationales related to blocked permissions, with labels like *"settings"* or *"go to settings."* Occasionally, this button can also serve a neutral function, dismissing the rationale without further steps. It is often labeled with phrases like *"ok"* or *"got it."*

Additionally, when a negative button is present alongside the positive button (46%^{s}), it is utilized to prevent the permission request from emerging. This button can also manifest as a dismiss *"x"* button located in the upper right corner of the rationale (4%^{s}). The negative button is often labeled with phrases like *"cancel," "deny," "don't allow," "later," "not now," "skip,"* or *"close."* In rare cases, a button providing an alternative solution might replace the negative button, like a button to manually enter the current location (2%^{s}).

App developers might employ an opinionated design for buttons to motivate users to grant permission. When there are multiple buttons, one approach is to make the positive button stand out more prominently than the negative button (19%^{s}).

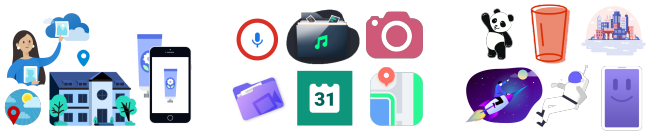

(a) Feature-related (b) Permission-related (c) Unrelated

Fig. 9: Icon & image types in rationales.

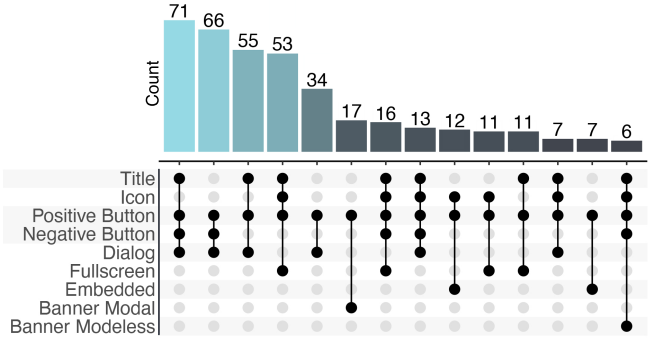

Fig. 10: Most common design element combinations (dataset: screenshots $N = 428$).

Images & Icons. Rationales can include images and icons (33%^{s}). We discovered that these visual elements are often (89%^{s}) directly related to the protected functionality (see Figure 9a) or the requested permission (see Figure 9b). In other instances (11%^{s}), these images and icons might be the app's logo or purely for visual appeal (see Figure 9c).

Title. The presence of a title (61%^{s}) depends on the rationale's layout. Generally, embedded rationales, modeless banners, or concise dialogs do not include titles. In most cases, the title aligns with the content of the rationale message, using phrases like *"permission required"* or *"grant permission"* (54%^{s}). Alternatively, a more neutral wording could be *"permission request"* or *"location permission"* (26%^{s}). Occasionally, the title serves as a welcome message to the app or states the app's name (8%^{s}). Titles can also serve as attention-grabbers; we observed that some titles incorporated an exclamation mark icon, prompting users to take action, as seen in examples like *"Warning!," "Attention!,"* and *"Action Required!"* (12%^{s}).

Summarizing the rationale designs in Figure 10, we observe distinct patterns across different layout types. Dialogs are the most common layout, typically including a positive button, a negative button, and often a title. Fullscreens, taking advantage of additional space, can include extra information, visual elements, a title, and usually a positive button only. Embedded rationales often incorporate an icon or image along with a positive button. On the other hand, modeless banners contain only a positive button. Modal banners can include a title, a visual element, and both positive and negative buttons.

Finding: Our investigation of rationales in Android apps showed considerable variation in how developers implement rationales in terms of phrasing and design. Nevertheless, we also identified some common trends that developers followed more frequently. Many developers preferred using dialogs to present rationales. Furthermore, we observed that rationales tend to be concise, typically composed of a single phrase from one perspective—user, app, or objective.

IV. User Study

Our exploratory analysis in revealed a variety of patterns in how app developers present rationales. Building on this insight, we conduct a user study to gather direct user feedback on various rationale phrasings. Our study evaluates users' permission choices, their understanding of these decisions, satisfaction levels, and perceived control. Ultimately, we aim to establish practical guidelines that aid app developers in effectively communicating permission requirements to enhance the efficacy of rationales and improve overall user experience.

A. Study Design

We structured the study as an online experiment using a repeated measures (within-subject) design. This approach was chosen to minimize errors related to individual differences, which are often misrepresented in between-subjects designs in judgment-related studies [7]. Each participant interacted with rationales for the four most common permissions—location, storage, camera, and microphone—as discussed in Section III-A1. These rationales were constructed by combining different rationale building blocks derived from our exploratory analysis, detailed below. To account for potential similarities among observations from the same user, we implemented a multilevel design, illustrated in Figure 11. Additionally, we randomized the sequence of rationales and permissions to mitigate order effects.

We maintained a consistent rationale design throughout our user study. This approach acknowledges that testing all conceivable designs and formulations in a single user study would be economically and statistically infeasible. Therefore, our strategy involves identifying the most effective phrasing for rationales, which can then be tested across diverse designs in future research. This strategy mirrors recent research that separately examines phrasing and design patterns [29]. Additionally, the appearance of rationales is often tied to a specific application or platform and cannot be uniformly governed. For the user study, we used dialogs because they are most commonly used for rationales, look similar across different apps, and are suitable for sharing important information that needs quick attention, like granting or denying permission.

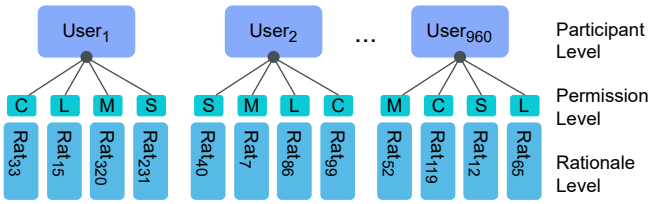

Fig. 11: Hierarchical structure of the user study. Permissions: C=Camera, L=Location, M=Microphone, S=Storage.

Rationales can accompany permissions requested either upfront (e.g., at app launch) or in context (e.g., button click). Previous research has demonstrated that rationales overpower the effect of timing [19], prompting our focus on upfront rationales. Additionally, we focused on type I rationales identified in our exploratory analysis, which are presented before permission is requested and visible to all users. In contrast, type II and III rationales are shown exclusively to users who deny permission, potentially multiple times.

We would like to emphasize that our study was preregistered to enhance transparency and credibility. Preregistration enabled us to define our objectives, sample size considerations, and study models upfront, minimizing biases and ensuring robust findings. For more details, please see our preregistration on the Open Science Framework [17].

1) Rationale Building Blocks for the User Study: The linguistic features from manually coded real-world rationales informed the rationale building blocks for our user study. Each rationale in the study includes a fundamental building block and may include one optional building block. This limitation restricts the number of phrases in each rationale, aiming to balance the systematic testing of various rationale combinations with participants' cognitive processing limits [14]. This design choice also mirrors the brevity observed in rationales from our exploratory analysis (see Section III). Table I presents the phrasings for each rationale building block utilized in the study. When multiple options were available for a block, we prioritized the most common choice. For instance, for the [user-demanding] block, we selected *"you need"* over less common options like *"you must."* Depending on the building blocks, each rationale is tied to a specific permission, adopts a certain perspective, includes clear or vague functionality, is phrased either positively or negatively, and can incorporate additional information. Furthermore, the functionality can also be user-centered, such as *"to scan your documents,"* which is paired with rationales phrased from the users' perspective.

Combining the rationale building blocks resulted in a vignette experiment with five dimensions (see Table I), with 320 unique vignettes/rationales ($4 \times 5 \times 2 \times 2 \times 4$), 80 per permission type. We presented each participant with four

TABLE I: Rationale building blocks for the user study.

Permission (×4):

[camera]	to scan (your) documents.
[location]	to search for gas stations (near you).
[microphone]	to send voice messages to (your) contacts.
[storage]	to attach photos to (your) posts.

Perspective & Politeness (×5):

[user-demanding]	You need to allow...
[user-polite]	Please allow...
[app-demanding]	We need...
[app-polite]	We would like...
[objective]	Camera permission is needed...

Functionality (×2):

[specific]	...to scan (your) documents.
[vague]	...(for the app) to work correctly.

Articulation (×2):

[positive]	...to scan (your) documents.
[negative]	Without this permission, you cannot...

Additional Information (×4):

[none] _
[guarantee] We do not collect or transfer any personal data outside your phone.
[control] You can change permissions from device settings anytime.
[more] For more information, see the privacy policy on our website.

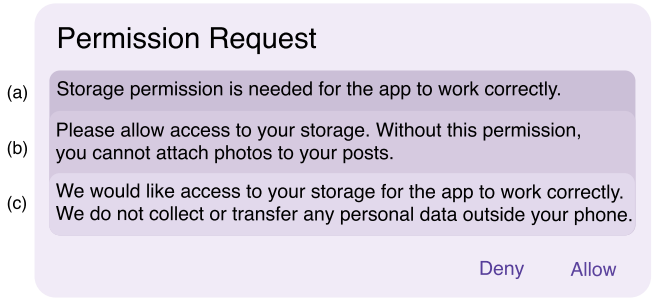

Fig. 12: Example of three distinct rationale sentences consisting of different building block combinations. (a) objective, vague, positive, none, (b) user-polite, specific, negative, none, (c) app-polite, vague, positive, guarantee.

randomly selected rationales, each on a different permission, in random order. We also ensured an even distribution of participants among the five dimensions. Figure 12 shows example rationales for three different vignettes.

B. Procedure

The study was conducted as an online experiment via the survey software Qualtrics. Upon granting their consent, participants were provided with a brief introduction to the study procedure. The central part of the study was then carried out and repeated four times. In each iteration, participants encountered a randomly selected rationale from a pool of 80 rationales per permission. During this phase, participants were instructed to carefully read the presented rationale and indicate their decision to grant or deny the corresponding permission. Afterward, participants were reminded of their previous decision on a separate screen. They were asked to evaluate their decision. Lastly, participants were requested to provide demographic information. The study procedure and measurements were refined based on insights from a pilot study involving 5 participants. To learn more about the study procedure, please consult Appendix B, and for further details on measurements, refer to Section IV-D.

C. Recruitment and Incentives

Participants were recruited using Prolific [56], ensuring a balanced sample of male and female participants. Each participant received £2 for completing the 10-minute survey (£12.00/hour). To be eligible, participants needed to be at least 18 years old, fluent in English, and regular users of mobile phones. We included participants from different mobile operating systems, like Android and iOS, as linguistic features can influence any user regardless of their OS.

To determine the optimal sample size for our study, we employed Monte Carlo simulations of the relevant multilevel models [37]. Without effect size estimates derived from prior research or meta-analyses, we assumed a standardized effect size of 0.25, which falls within the range of a medium to large effect size [24], [52]. This choice of effect size is particularly appropriate as the study's objective is to derive practical recommendations for app developers, necessitating a difference that holds practical relevance. Additionally, our sample size's boundaries were guided by increasing the sample in multiples of full presentations of all available rationales (e.g., 320 $\rightarrow$ 640 $\rightarrow$ 960) to achieve an equal distribution between rationales. Given these boundaries, our simulations indicated that a sample size of 960 participants could detect a fixed effect of 0.25 with a power of at least 0.8 while maintaining an alpha error probability of 0.05.

In total, we gathered data from 980 participants. To maintain data quality, we followed recommendations from the literature for filtering out careless responses [68]. We excluded 18 individuals who self-reported to not use their data, one person due to missing data, and another person because they completed the task three times faster than the median speed of the sample. Our final sample consisted of 960 participants, 49.8% of whom identified as female and 48.5% as male. Additionally, 16 participants identified with other self-reported genders or chose not to disclose. Participants had a mean self-reported age of 32 ($SD = 10.8$ years). Most participants attended college, with 50.3% earning an undergraduate degree, 17.9% holding a graduate degree, and 17.1% not completing their studies. Regarding smartphones, 60.3% used Android, 38.1% used an iPhone with iOS, and 1.6% used a Windows phone, which is comparable to the worldwide mobile OS marketshare [28].

To ensure our study represents a worldwide viewpoint and to increase the relevance of our findings across cultural boundaries, we recruited participants from different countries as long as they spoke fluent English. Consequently, our final international sample included participants from a wide range of geographic regions spanning multiple continents, as facilitated by Prolific (45% from Europe, 37% from the Americas, 13% from Africa, 4% from Oceania, and 1% from Asia), representing a total of 31 different countries. Refer to Appendix A for a detailed breakdown of participant countries.

D. Measurements

In our study, we utilized a range of measurements, which are explained in this section and are also available in the questionnaire provided in Appendix B.

1) Decision & Decision Evaluation: Participant's decision was assessed with one item asking them to choose between *"allow"* or *"deny"*, which resembled the choice in the presented app screenshot of the rationales.

To assess users' evaluation of their decisions, we used an adapted version of the Decision Evaluation Scales (DES), which was successfully used in this context before [19]. The scale was initially adopted from the field of health psychology [62], where it was designed to evaluate patients' decision to uptake or refuse a treatment choice, which is analogous to users' decision to grant or deny a permission request. Additionally, the DES allows the investigation of multiple essential dimensions of the decision evaluation, such as (1) whether users received sufficient information to make an informed decision, (2) their satisfaction with the decision, and (3) their perceived control over the decision.

For the user study, we extended the scale used in the related work by adding one additional item to each subscale. In this way, we aimed to increase the reliability of the assessment [2], capture a broader perspective on the underlying construct, and prevent censored scale averages due to low item difficulty. To create additional items for the scale, we constructed a set of five new items for each subscale, which were then subsequently rated by a sample of nine domain experts (four behavioral scientists and five information security experts). The three items with the highest agreement were then added to the user study. The scale items, the results for multilevel internal consistency, and factor loadings for all items from confirmatory factor analysis are given in Appendix C. All the subscales of the DES were measured on a seven-point rating scale, with scores closer to 7 indicating greater agreement with the item and scores closer to 1 indicating greater disagreement.

2) Person Level Measurements: To account for individual variations, besides gathering demographic data (like gender, age, educational level, and users' mobile OS), we also assessed participants' privacy concerns and their past privacy experiences. Previous research has indicated that these factors can influence users' decisions regarding runtime permission requests [9], [19], [72].

E. Ethical Considerations

The study was approved by our institution's Ethics Review Board. Data collected via Prolific and Qualtrics were treated sensitively, with personal identifiers separated. At the start of the study, participants received precise details about the purpose of the study and the data being collected. We ensured participants understood how their data would be used while allowing them to withdraw their participation at any time.

F. Model Construction

We used linear multilevel models to test whether rationale phrasings impacted users' decision-making process and conducted all analyses in R version 4.2.2 (R Core Team, 2023). As part of our data preparation, we computed scale means for measurements with multiple items, including informed decision, decision satisfaction, decision control, prior privacy experience, and privacy concerns. Additionally, we standardized all user-level predictors (age, prior privacy experience, and privacy concerns) by using grand mean centering. Categorical predictors were coded using treatment coding, with the reference groups as follows: perspective & politeness (objective), articulation (negative), functionality (vague), and additional information (none). We deviated from our preregistration for permissions and chose to treat them as random effects rather than fixed effects with difference coding. This change did not substantially affect the impact of other variables but allowed us to investigate the variation between permissions more closely.

We took a step-by-step approach to simplify our modeling process and ensure consistency with recommendations from prior research [38]. We used maximum likelihood estimation for all models to make them comparable. We built and tested the models as follows: (1) In the first step, we started with a simple regression model. Next, we expanded it to a random intercept model, considering permission and user as random effects. (2) For the second step, we introduced control variables, specifically prior privacy experience and privacy concerns. For the DES models, we also incorporated participants' decisions as a control variable because the outcome (i.e., granting or denying a permission request) could affect users' comfort level with their choices. (3) Progressing to the third step, we added the variables of interest. These included perspective & politeness, articulation, functionality, and additional information. In the fourth step (4) we tested adding interactions between the variables of interest. However, including these did not enhance the model fit. For more details about the model-building process and fit criteria, see Appendix D. The final models were recalculated using Restricted Maximum Likelihood Estimation, which leads to a more conservative and less error-prone estimation of the parameters [38]. Table II shows the final model for each outcome variable.

G. Results

We now present the results of our user study, which we discuss and interpret subsequently in Section V.

1) Effect of Rationale Building Blocks: Below are the outcomes of the effects of the five rationale building blocks.

Perspective & Politeness: When rationales were phrased politely from the app's perspective (e.g., *"we would like"*), it had an interesting effect. The likelihood of granting permissions decreased ($odds\ ratio = 0.71$, $std.\ \beta = 0.71$, $p = 0.003$). However, whether the rationale was phrased from different perspectives and had a polite or demanding tone did not influence how participants perceived their decision.

Functionality. Rationales that explained why permission is needed had a positive impact on participants, making them feel more informed about their decision ($\beta = 0.16$, $std.\ \beta = 0.11$, $p < 0.001$) and more satisfied with their choices ($\beta = 0.07$, $std.\ \beta = 0.07$, $p = 0.018$). However, this did not

TABLE II: The final multilevel models.

	Decision *Odds Ratio* (*std.* β)	DES Inform β (*std.* β)	DES Satis β (*std.* β)	DES Control β (*std.* β)
(Intercept)	1.76 (1.76)***	-1.01 (-0.61)***	1.68 (0.19)***	0.27 (-0.01)***
Privacy Concerns	0.71 (0.71)***	0.12 (0.11)***	0.09 (0.11)***	0.03 (0.03)
Prior Experience	0.89 (0.88)**	-0.02 (-0.02)	-0.08 (-0.08)***	-0.17 (-0.16)***
Decision Grant	–	1.08 (0.75)***	-0.30 (-0.28)***	-0.15 (-0.12)***
Perspective & Politeness (ref: objective)				
user-demanding	0.87 (0.87)	0.05 (0.03)	0.00 (0.00)	-0.03 (-0.02)
user-polite	0.86 (0.86)	0.01 (0.01)	0.04 (0.04)	0.00 (0.00)
app-demanding	0.80 (0.80)	-0.00 (-0.00)	-0.04 (-0.04)	0.02 (0.02)
app-polite	0.71 (0.71)**	0.06 (0.04)	0.00 (0.00)	0.07 (0.06)
Functionality (ref: vague)				
specific	1.12 (1.12)	0.16 (0.11)***	0.07 (0.07)*	0.03 (0.03)
Articulation (ref: negative)				
positive	1.03 (1.03)	-0.02 (-0.01)	-0.09 (-0.08)**	-0.02 (-0.02)
Additional Information (ref: none)				
guarantee	1.43 (1.43)***	0.24 (0.16)***	0.06 (0.05)	0.17 (0.15)***
control	1.28 (1.28)*	0.06 (0.04)	-0.01 (-0.01)	0.13 (0.11)**
more	1.07 (1.07)	0.11 (0.08)*	-0.07 (-0.07)	0.05 (0.05)
Marginal R^2	0.050	0.141	0.037	0.029
Conditional R^2	0.177	0.414	0.381	0.499

Models were fitted with the Restricted Maximum Likelihood estimation. The coefficients for Decision are shown as odds ratios, where values <1 indicate that the likelihood of granting permissions is lower than the likelihood of denying the permission and values >1 indicate that the likelihood of granting the permission is higher. *std.* β = *standardized* β. $^{*}p < .05$, $^{**}p < .001$, $^{***}p < .0001$. Decision coding: 0 = *deny*, 1 = *allow*. N: 960_{User}, $3840_{\text{Rationale}}$.

significantly increase the likelihood of granting permissions or make participants feel more in control of their decisions.

Articulation. The tone used in the rationale also played a role. Positively phrasing the rationale, highlighting the benefits of granting permission, decreased decision satisfaction ($\beta = -0.09,\ std.\ \beta = -0.08,\ p = 0.005$). However, this positive phrasing did not affect the decision itself, the perception of being informed, or the sense of control over the decision.

Additional Information (Guarantee). Guaranteeing that the permission will not be misused had a notable impact. It significantly increased the likelihood of participants granting permission ($odds\ ratio = 1.43,\ std.\ \beta = 1.43,\ p < 0.001$). Furthermore, participants felt more informed about their decision ($\beta = 0.24,\ std.\ \beta = 0.16,\ p < 0.001$) and more in control ($\beta = 0.17,\ std.\ \beta = 0.15,\ p < 0.001$). Still, it did not have a direct effect on decision satisfaction.

Additional Information (Control). Informing participants that they can change their decision at any time increased the likelihood of granting permission ($oddsratio = 1.28, std.\beta = 1.28, p = 0.018$). It also made participants feel more in control of their decisions ($\beta = 0.13,\ std.\ \beta = 0.11,\ p = 0.002$). However, it did not impact participants' perception of their decision being informed or their satisfaction with the decision.

Additional Information (More). Interestingly, adding the phrase *"For more information, see the privacy policy on our website"* did not provide additional information but increased participants' perception of being informed ($\beta = 0.11, std.\beta = 0.08,\ p = 0.047$). Other variables, such as the decision itself, satisfaction, and control, remained unaffected.

2) Effect of the Generic Rationale: Our design included a generic base rationale, which consisted of a single negatively articulated phrase from the objective point of view. This phrase described a vague functionality, was neither polite nor included any additional building blocks (e.g., *"Without storage permission, the app cannot work correctly."*). The effect of this rationale can be investigated by interpreting the intercepts of our dependent variables. In this context, the baseline of value zero signifies either denying the permission request (Decision) or corresponds to the average score of four for the three judgments related to the rationales (DES Inform, DES Satis, DES Control). On all four variables, our results showed that participants deviated significantly from this baseline in terms of declining the permission request or feeling indifferent when presented with the generic rationale. Even with the generic phrasing, participants were more likely to grant the permission ($odds\ ratio = 1.76,\ std.\ \beta = 1.76,\ p < 0.001$). They felt less informed ($\beta = -1.01, std.\beta = -0.61, p < 0.001$), but still in control ($\beta = 0.27,\ std.\ \beta = -0.01,\ p < 0.001$) and satisfied ($\beta = 1.68,\ std.\ \beta = 0.19,\ p < 0.001$) with their decision.

3) Effect of Other Variables: We also examined the effects of other variables related to individual differences or known to influence permission decisions from prior research.

Privacy Concerns. Participants' privacy concerns had a multifaceted impact on their permission decisions and perceptions. Firstly, higher privacy concerns were associated with a decreased likelihood of granting permissions, indicating that individuals with privacy concerns were less inclined to provide access ($odds\ ratio = 0.77,\ std.\ \beta = 0.71,\ p < 0.001$). Conversely, higher privacy concerns had a positive influence

on participants' perception of making an informed decision ($\beta = 0.12$, $std.\,\beta = 0.11$, $p < 0.001$) and their satisfaction with the decision ($\beta = 0.09$, $std.\,\beta = 0.11$, $p < 0.001$). However, privacy concerns did not significantly impact participants' perception of control over their decisions.

Prior Privacy Experience. Examining participants' previous encounters with privacy-related experiences provided the following insights. Participants with more prior interactions with privacy issues were less likely to grant permissions ($oddsratio = 0.89$, $std.\beta = 0.88$, $p = 0.006$). This trend also extended to participants reporting lower overall satisfaction with their decisions ($\beta = -0.08$, $std.\,\beta = -0.08$, $p = 0.001$) and a diminished sense of control over their choices ($\beta = -0.17$, $std.\,\beta = -0.16$, $p < 0.001$). Prior privacy experience did not affect the perception of making an informed decision.

Permission Decision. Analyzing participants' decisions revealed that granting permission increased their sense of being informed ($\beta = 1.08$, $std.\,\beta = 0.75$, $p < 0.001$) but came with trade-offs. It lowered both decision satisfaction ($\beta = -0.30$, $std.\,\beta = -0.28$, $p < 0.001$) and the sense of control ($\beta = -0.15$, $std.\,\beta = -0.12$, $p < 0.001$). In essence, while granting permission made participants feel informed, it reduced satisfaction and control.

Age, Gender, Educational Background & Mobile OS. Previous studies in the field have investigated age, gender, educational background, and mobile OS as possible confounding variables [9], [11]. To investigate the effects of these variables on our main outcomes, we compared the models with these added control variables to the models reported above. For the control variables, we excluded answers with only a few observations (e.g., the "other" category in educational background). We refitted the models reported above on the reduced dataset ($N_{reduced} = 929$) to allow a numeric comparison depending on the common fit measures. Inspecting these comparisons, only the models for decision ($\chi^2(5,\, N = 929) = 12.72$, $p = 0.029$) and informed decision ($\chi^2(5, N = 929) = 12.524, p = 0.028$) improved significantly over the study models.

For these two models, none of the independent variables' effects were decisively affected by adding the additional control variables to the model. For decision, we found that participant's odds of granting permission were significantly increased if they held an undergraduate degree compared to a graduate degree ($oddsratio = 1.32$, $std.\beta = 1.32$, $p = 0.019$) and decreased if they were using iOS compared to Android ($odds\ ratio = 0.79$, $std.\,\beta = 0.79$, $p = 0.010$). Additionally, participants who identified as female felt they had made a slightly less informed decision ($\beta = -0.15$, $std.\,\beta = -0.11$, $p = 0.008$) compared to participants who identified as male. We did not find a significant effect of age.

4) Variation Between Clusters: A considerable proportion of the variance in the models was explained by differences within participants and permissions instead of only the fixed factors. This is expected as evaluating the decision to grant permission is a complex cognitive process most likely influenced by many different characteristics of the participants (e.g., the propensity to make a decision) and the provider of the rationale (e.g., the app's trustworthiness). Although we did not measure specific characteristics of the participants in this regard, our hierarchical analysis accounted for these differences in clusters of participants and permissions and uncovered that the base effect of the rationales mainly varies within participants, ranging from an intercept variance of $\tau_{00\_Satisfaction} = 0.40$ to $\tau_{00\_Control} = 0.64$, while there was only a small variance based on the permissions ($\tau_{00} = [0.00; 0.04]$).

V. Discussion

App developers have many choices when crafting rationales, including prompt styles and different phrasings. Despite the availability of various guidelines on creating rationales [5], [34], [53], our empirical analysis revealed significant diversity in developers' approaches. Additionally, studies in other domains have shown how linguistic variations in phrasing can affect various user decisions [1], [16], [29], [44], [60], [66]. Given these factors, it is essential to understand how the diverse phrasing of rationales can influence users' permission decisions and, consequently, their privacy. Focusing on existing guidelines, this discussion will show their adoption in our dataset, compare available recommendations with our findings, and distinguish our research from other studies on the trustworthiness of rationales. We will also provide directions for implementing our recommendations in future research.

A. Do Real-World Rationales Follow Guidelines?

Guidelines from Android [34], iOS [5], and NN/g [53] recommend the following for phrasing rationales: provide specific functionality, avoid passive voice, and include user-related features. In our dataset, we found that the majority of rationales adhere to these recommendations. For instance, more than half of the rationales we inspected explain why permission is necessary, specifying a particular functionality (58%^{m} 69%^{s}). Furthermore, many app developers avoid using passive voice (objective perspective in our analysis) in their rationales (73%^{m} 70%^{s}). However, only a small fraction of rationales emphasize the benefits to the user when granting permission (18%^{m} 17%^{s}), such as *"so you can make video calls,"* or *"to share pictures with your friends."* Even though there is relatively high compliance on the side of developers, many real-world examples phrase their rationales differently or opt to include additional details.

B. Influential Building Blocks Within and Beyond Guidelines

Every user possesses distinct privacy concerns and preferences. We tried to apprehend these nuances by capturing how users perceive and assess their permission choices, regardless of whether they grant or deny permissions. Our assertion is that users should feel well-informed, satisfied, and in control of their decisions, aligning with their individual privacy preferences within a given context. In summary, we provide actionable recommendations for developers based on the insights from the user study. These recommendations can serve as a more granular extension of the existing guidelines.

Provide specific functionality and phrase it negatively. In line with available guidelines, we found that clearly explaining why permission is required by specifying a functionality improves users' perception of having made an informed decision and increases their satisfaction with that decision. Notably, users tend to be more satisfied with their decision when permission requirements are presented in a negative context, such as "Without this permission, {certain app functionality} cannot be used." This negative phrasing appears to be more straightforward for users, enhancing their satisfaction with the decision-making process compared to a positive phrasing, such as "This permission is used for {certain app functionality}."

We also found that supplementary information blocks positively influenced users. However, given the limited space available for rationale messages, we recommend that app developers prioritize these building blocks in order of their effectiveness, as space permits.

Assure users what the permission will not be used for. The first and most influential addition is including a guarantee that permission will not be misused. Our study found that this addition significantly enhances users' perception of making an informed decision, empowers them to feel more in control, and increases grant rates. Users were influenced by such assurances even without concrete proof, possibly due to perceived transparency from the app developers or a misbelief that app distribution platforms prevent fraudulent permission requests. As such misconceptions pose a risk to users, it is essential that developers include this building block only if the app genuinely upholds these promises and provides a legally binding privacy policy. Furthermore, we recommend that the truthfulness of this statement should be verified using information from the rationale during app vetting. In fact, the mandated presence of such useful information supports the vetting process, as detailed in Section V-D.

Highlight the reversibility of permission decisions. The second addition we examined is reminding users that they can modify their permission choices, articulated as "You can change permissions from device settings anytime." This phrase enhanced users' perception of control and increased their likelihood of granting permission. This effect aligns with the control paradox, which posits that the perception of control increases the likelihood of disclosing sensitive information [10]. It may also stem from the notion that users feel less apprehensive when they have a sense of control over their decisions, similar to being the driver of a car rather than a passenger. This finding is consistent with prior research, which shows that users are more likely to grant permission when aware of the option to change their decision later [9]. However, other studies indicate that users rarely exercise this option [49], [61]. Therefore, highlighting the reversibility of permission decisions could increase user awareness and positively influence their perception of the decision. However, adding this information to a rationale is not enough; we must also translate this sense of control into actionable steps. Given the benefits of knowing that decisions can be revoked, future research could focus on effectuating this sense of control, such as nudging users to review their previous permission choices [3], [72] or providing more fine-grained permission controls, akin to one-time granting of some permissions on Android [33] and iOS [4].

Provide supplemental information. The third addition is the phrase "For more information, see the privacy policy on our website." Although this phrase does not provide specific information about the purpose of the permission, it enhances users' perception of being informed. This outcome may result from the impression of users that the app developer values transparency, offers supplementary information, and complies with legal requirements. These factors collectively foster a sense of being well-informed among users, despite the absence of explicit details or direct links to additional material in the rationale. However, it is possible that the users in our study, who were prompted to imagine a real situation, just acted on the assumption that in a real-life situation, they would have been able to obtain more information about the app and its features. If not, this finding is somewhat worrisome because no further informational value was added to the rationale, but still, users felt better informed. Nevertheless, in practice, we recommend adding a link to supplemental information for users who need more information to decide. Going a step further, our results indicate that users may assume any additional information, even just a hint to a privacy policy, is verified and trustworthy. It appears that users delegate the vetting process, to other users or Google. Therefore, we see an opportunity for future research to shift this implicit user trust to explicit verification. Ideally, rationales should distinguish between provided and verified privacy policies, highlighting and differentiating them.

C. Revisiting Rationale Guidelines

We found that not all available guidelines influenced our outcome variables as expected. For example, the specificity of the provided functionality had no significant effect on users' permission decisions. Whether specifying the use of permissions has no impact or only a minor effect on grant rates remains undetermined, as our study design may not have detected these nuances. However, our observation aligns with previous research, which has shown that any explanation provided by developers tends to increase grant rates, regardless of the explanation's meaningfulness [65].

Additionally, contrary to guidelines advising against passive voice, we found that using the passive voice (the objective perspective) in rationales did not have a detrimental impact. In fact, in some cases, rationales written in the passive voice had a higher likelihood of being granted compared to those phrased politely from the app's perspective, such as "we would like your permission." Although not directly examined in our study, we suspect that polite language might make users perceive the permission as optional, resulting in lower grant rates. Moreover, we did not find a significant effect of adding user-related features to the rationale.

In conclusion, this does not mean that these guidelines do not contribute to the overall readability or user comprehension;

it simply means that we did not find a significant impact on our main outcome variables, which we consider relevant indicators of improving the overall user experience. Thus, our results suggest that existing guidelines can still be enhanced, for example, by adding supplementary information. Additionally, we found that the perspective of a rationale is less critical than previously thought. Finding further gaps might also be an interesting question for future research.

D. Differentiating Between Usability and Trustworthiness

In this work, our efforts prioritize enhancing the usability of rationales, which is distinct from addressing their trustworthiness and verification. This differentiation mirrors similar distinctions in research concerning other self-reported privacy instruments such as privacy policies, Privacy Nutrition Labels, and Google's Data Safety Section. While some studies concentrate on enhancing the usable security aspects [15], [48], [74], [75], others focus on improving and assessing the trustworthiness of these instruments [26], [40], [42], [43], [45], [64], [71]. Due to this distinction, it is important to note that malicious app developers could potentially exploit the above phrasing strategies, which enhance the usability of app rationales, to the user's disadvantage. Consequently, ensuring the trustworthiness of rationales remains an important question that needs to be addressed by parallel works. For example, app vetting can use rationales as reference points in analyzing and classifying app behavior. Ultimately, both aspects are indispensable: a usable yet unverified rationale holds no more value than an unusable verified rationale. If rationales are acknowledged as effective privacy declarations by end-users, it becomes imperative to mandate app developers to provide them in a structured format. This enhances user comprehension and allows for the development of solutions that can analyze privacy compliance on a large scale, much like privacy nutrition labels [75].

E. Future Directions for Guideline Adoption

Having discussed the available guidelines and provided more nuanced recommendations, the question of how to best assist developers in implementing these guidelines remains open. On one hand, we believe that rationale guidelines should remain recommendations to guide app developers toward best practices without being enforced, allowing developers the freedom to create rationales aligned with their corporate identity. On the other hand, given that the guidelines proposed in this paper emphasize the phrasing of rationales, there is significant potential for developing a tool for modular rationales. Such a tool would allow developers to utilize rationale building blocks, enabling them to assemble customized rationales that suit their specific needs. This approach could reduce the burden on developers in crafting rationales while still giving them the flexibility to design the UI of rationales that reflect their corporate identity. This tool could also be used for iOS rationales which, in specific cases, can be provided in custom pre-alert screens [5] in addition to a purpose string that is integrated into the permission request dialog.

VI. Threats to Validity

Our user study is subject to methodical deliberations, which constrain its scope in terms of certain forms of validity. In pursuit of high internal validity, we implemented a rigorous experimental design that held all influencing factors constant, thus increasing the likelihood that the results reflect only the actual effects of rationale building blocks. However, this approach required presenting permission requests through vignettes rather than asking users to install an app on their device that monitors the handling of permission requests and subsequently poses questions. Limited external validity is inherent in every vignette study, as it prompts users to imagine the situation rather than being in it, thus not fully reflecting real-world permission request handling. However, continuously monitoring user's device usage also raises significant ethical concerns. In our case, no prior studies existed that would justify infringing on a user's privacy without knowing if rationale phrasing would have a measurable effect on their decisions. Instead, our approach allowed us to test numerous possible variations under controlled conditions, whereas a field study would rely on a sample of rationales from incidental app installations, which also limits its generalizability. Nevertheless, we still displayed the rationale in the design of a genuine app prompt on a smartphone to enhance immersion and realism. Consequently, we contend that our chosen study design was the most appropriate under the given circumstances.

Additionally, our use of a repeated measurement design may have caused participants to suspect the research purpose to some extent when encountering the second rationale. To mitigate potential sequence effects, we employed randomization by assigning random permission orders for each participant and presenting them with a randomly selected set of rationales. Additionally, we explicitly instructed users to assume that every prompt represented a new app installation. While these measures counteract common sequence effects, they cannot completely eliminate the possibility of user boredom or careless responses due to repetition. Therefore, we conducted rigorous data analysis to identify indications of such issues and handled them appropriately.

Furthermore, it is important to note that we did not explore the impact of different rationale designs, which could potentially have a significant effect on user perception. Our choice to focus on a single design was primarily driven by feasibility considerations, as discussed in Section IV-A. Investigating multiple design parameters alongside the rationale variations would also have complicated our statistical models and might have led to issues with multicollinearity. Therefore, we opted for a consistent rationale design that was commonly observed in the rationales we examined. Nonetheless, we encourage further research on this topic, building upon the comprehensive investigation of design variations reported in this study.

On a final note, while the results of our study demonstrate the influence of rationale variations on users, it is important to recognize that they have not been field-tested, where other contextual factors may play a significant role. To accurately

anticipate the real-world impact of our recommendations, it is essential to remember that human behavior is rarely influenced by a single, straightforward cause, and rationales are just one factor among many. Thus, when applying and interpreting these results, it is crucial to keep in mind that we are dealing with subtle nuances in language and content rather than dramatic changes in design.

VII. Conclusion

In this work, we extensively examined real-world rationales in the context of mobile app permissions. Through this investigation, we uncovered diverse building blocks and design elements of rationales. Our user study, involving 960 participants, unveiled the impact of phrasing on users' permission decisions and their assessment of those decisions. By aligning our findings with established recommendations and guidelines, we extracted valuable insights, offering actionable recommendations for app developers to enhance user experience through more effective rationale crafting. Our work underscores the importance of well-considered phrasing of rationales and extends an invitation for future research. Subsequent investigations could explore additional dimensions, emphasizing refining the design and overall user experience of rationales. Furthermore, establishing rationales as a usable and standardized instrument can yield enhancements in other areas, such as app vetting, including the validation of rationale messages and Google's Data Safety Section.

References

[1] O. Akgul, R. Abu-Salma, W. Bai, E. M. Redmiles, M. L. Mazurek, and B. Ur, "From secure to military-grade: Exploring the effect of app descriptions on user perceptions of secure messaging," in *Proc. 20th Workshop on Privacy in the Electronic Society (WPES'21)*, 2021.

[2] M. J. Allen and W. M. Yen, *Introduction to Measurement Theory*. Waveland Press, 2002.

[3] H. Almuhimedi, F. Schaub, N. M. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, "Your location has been shared 5,398 times!: A field study on mobile app privacy nudging," in *Conference on Human Factors in Computing Systems (CHI'15)*, 2015.

[4] Apple, "About privacy and location services on iOS, ipadOS and watchOS," https://support.apple.com/en-gb/102515, accessed: 2024-07-01.

[5] ——, "Privacy," https://developer.apple.com/design/human-interface-guidelines/privacy, accessed: 2024-07-01.

[6] ——, "Transparency is the best policy," https://www.apple.com/privacy/labels, accessed: 2024-07-01.

[7] M. Birnbaum, "How to show that $9 > 221$: Collect judgments in a between-subjects design," *Psychological Methods*, vol. 4, pp. 243–249, 1999.

[8] K. Bongard-Blanchy, J. Sterckx, A. Rossi, V. Distler, S. Rivas, and V. Koenig, "An (un)necessary evil - users' (un)certainty about smartphone app permissions and implications for privacy engineering," in *Proc. 7th IEEE European Symposium on Security and Privacy (EuroS&P'22)*, 2022.

[9] B. Bonné, S. T. Peddinti, I. Bilogrevic, and N. Taft, "Exploring decision making with android's runtime permission dialogs using in-context surveys," in *Proc. 13th Symposium on Usable Privacy and Security (SOUPS'17)*, 2017.

[10] L. Brandimarte, A. Acquisti, and G. Loewenstein, "Misplaced confidences: Privacy and the control paradox," *Social Psychological and Personality Science*, vol. 4, pp. 340–347, 2013.

[11] W. Cao, C. Xia, S. T. Peddinti, D. Lie, N. Taft, and L. M. Austin, "A large scale study of user behavior, expectations and engagement with android permissions," in *Proc. 30th USENIX Security Symposium (SEC'21)*, 2021.

[12] S. Chen, L. Fan, C. Chen, and Y. Liu, "Automatically distilling storyboard with rich features for android apps," *CoRR*, vol. abs/2203.06420, 2022.

[13] S. Chen, L. Fan, C. Chen, T. Su, W. Li, Y. Liu, and L. Xu, "Storydroid: automated generation of storyboard for android apps," in *Proc. 41th International Conference on Software Engineering (ICSE'19)*, 2019.

[14] N. Cowan, "The magical number 4 in short-term memory: A reconsideration of mental storage capacity," *Behavioral and Brain Sciences*, vol. 24, pp. 87—-114, 2001.

[15] L. F. Cranor, "Mobile-app privacy nutrition labels missing key ingredients for success," *Commun. ACM*, vol. 65, pp. 26–28, 2022.

[16] V. Distler, T. Gutfleisch, C. Lallemand, G. Lenzini, and V. Koenig, "Complex, but in a good way? how to represent encryption to non-experts through text and visuals – evidence from expert co-creation and a vignette experiment," *Computers in Human Behavior Reports*, vol. 5, p. 100161, 2022.

[17] Y. Elbitar, A. Hart, and S. Bugiel, "Investigating the effect of rationales," https://osf.io/7zprb, accessed: 2024-07-01.

[18] ——, "Rationale codebook landscape," https://osf.io/z9huf, accessed: 2024-07-01.

[19] Y. Elbitar, M. Schilling, T. T. Nguyen, M. Backes, and S. Bugiel, "Explanation beats context: The effect of timing & rationales on users' runtime permission decisions," in *Proc. 30th USENIX Security Symposium (SEC'21)*, 2021.

[20] Explosion, "Prodigy," https://prodi.gy, accessed: 2024-07-01.

[21] ——, "SpaCy," https://spacy.io, accessed: 2024-07-01.

[22] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. A. Wagner, "Android permissions: user attention, comprehension, and behavior," in *Proc. 8th Symposium on Usable Privacy and Security (SOUPS'12)*, 2012.

[23] Y. Feng, L. Chen, A. Zheng, C. Gao, and Z. Zheng, "Ac-net: Assessing the consistency of description and permission in android apps," *IEEE Access*, vol. 7, pp. 57 829–57 842, 2019.

[24] D. C. Funder and D. J. Ozer, "Evaluating effect size in psychological research: Sense and nonsense," *Advances in Methods and Practices in Psychological Science*, vol. 2, pp. 156–168, 2019.

[25] H. Gao, C. Guo, Y. Wu, N. Dong, X. Hou, S. Xu, and J. Xu, "Autoper: Automatic recommender for runtime-permission in Android applications," in *Proc. 43rd IEEE Annual Computer Software and Applications Conference (COMPSAC'19)*, 2019.

[26] J. Gardner, Y. Feng, K. Reiman, Z. Lin, A. Jain, and N. Sadeh, "Helping mobile application developers create accurate privacy labels," in *Proc. 7th IEEE European Symposium on Security and Privacy (EuroS&P'22)*, 2022.

[27] G. J. Geldhof, K. J. Preacher, and M. J. Zyphur, "Reliability estimation in a multilevel confirmatory factor analysis framework," *Psychological methods*, vol. 19, pp. 72–91, 2014.

[28] GlobalStats, "Mobile operating system market share worldwide," https://gs.statcounter.com/os-market-share/mobile/worldwide, accessed: 2024-07-01.

[29] M. Golla, G. Ho, M. Lohmus, M. Pulluri, and E. M. Redmiles, "Driving 2fa adoption at scale: Optimizing two-factor authentication notification design patterns," in *Proc. 30th USENIX Security Symposium, (SEC'21)*, 2021.

[30] Google, "Activitycompat," https://developer.android.com/reference/androidx/core/app/ActivityCompat#shouldShowRequestPermissionRationale(android.app.Activity,java.lang.String), accessed: 2024-07-01.

[31] ——, "Dialogs," https://developer.android.com/develop/ui/views/components/dialogs, accessed: 2024-07-01.

[32] ——, "Localize your app," https://developer.android.com/guide/topics/resources/localization, accessed: 2024-07-01.

[33] ——, "One-time permissions," https://developer.android.com/training/permissions/requesting#one-time, accessed: 2024-07-01.

[34] ——, "Request runtime permissions," https://developer.android.com/training/permissions/requesting, accessed: 2024-07-01.

[35] ——, "Understand app privacy & security practices with google play's data safety section," https://support.google.com/googleplay/answer/11416267, accessed: 2024-07-01.

[36] ——, "Write automated tests with ui automator," https://developer.android.com/training/testing/other-components/ui-automator, accessed: 2024-07-01.

[37] R. Harrison, "Introduction to monte carlo simulation," *AIP conference proceedings*, vol. 1204, pp. 17–21, 2010.

[38] J. Hox, M. Moerbeek, and R. van de Schoot, *Multilevel Analysis: Techniques and Applications (3nd ed.)*. Routledge, 2017.
[39] iBotPeaches, "Apktool," https://apktool.org, accessed: 2024-07-01.
[40] A. Jain, D. Rodriguez, J. M. del Álamo, and N. M. Sadeh, "ATLAS: automatically detecting discrepancies between privacy policies and privacy labels," in *Proc. 8th IEEE European Symposium on Security and Privacy (EuroS&P'23)*, 2023.
[41] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. M. Sadeh, and D. Wetherall, "A conundrum of permissions: Installing applications on an android smartphone," in *Proc. 16th International Conference on Financial Cryptography and Data Security (FC'12)*, 2012.
[42] S. Koch, M. Wessels, B. Altpeter, M. Olvermann, and M. Johns, "Keeping privacy labels honest," *Proc. Priv. Enhancing Technol.*, vol. 2022, pp. 486–506, 2022.
[43] K. Kollnig, A. Shuba, M. V. Kleek, R. Binns, and N. Shadbolt, "Goodbye tracking? impact of ios app tracking transparency and privacy labels," in *Proc. 5th ACM Conference on Fairness, Accountability, and Transparency (FAccT'22)*, 2022.
[44] I. P. Levin, S. L. Schneider, and G. J. Gaeth, "All frames are not created equal: A typology and critical analysis of framing effects," *Organizational Behavior and Human Decision Processes*, vol. 76, pp. 149–188, 1998.
[45] T. Li, L. F. Cranor, Y. Agarwal, and J. I. Hong, "Matcha: An IDE plugin for creating accurate privacy nutrition labels," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 8, pp. 33:1–33:38, 2024.
[46] T. Li, K. Reiman, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding challenges for developers to create accurate privacy nutrition labels," in *Conference on Human Factors in Computing Systems (CHI'22)*, 2022.
[47] Y. Li, D. Chen, T. Li, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding ios privacy nutrition labels: An exploratory large-scale analysis of app store data," in *Conference on Human Factors in Computing Systems (CHI'22)*, 2022.
[48] Y. Lin, J. Juneja, E. Birrell, and L. F. Cranor, "Data safety vs. app privacy: Comparing the usability of android and ios privacy labels," *CoRR*, vol. abs/2312.03918, 2023.
[49] B. Liu, J. Lin, and N. M. Sadeh, "Reconciling mobile app privacy and usability on smartphones: could user privacy profiles help?" in *Proc. 23rd International World Wide Web Conference (WWW'14)*, 2014.
[50] X. Liu, Y. Leng, W. Yang, W. Wang, C. Zhai, and T. Xie, "A large-scale empirical study on Android runtime-permission rationale messages," in *IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC'18)*, 2018.
[51] X. Liu, Y. Leng, W. Yang, C. Zhai, and T. Xie, "Mining android app descriptions for permission requirements recommendation," in *26th IEEE International Requirements Engineering Conference (RE'18)*, 2018.
[52] J. Lorah, "Effect size measures for multilevel models: definition, interpretation, and timss example," *Large-scale Assessments in Education*, vol. 6, 2018.
[53] Nielsen Norman Group, "3 design considerations for effective mobile-app permission requests," https://www.nngroup.com/articles/permission-requests, accessed: 2024-07-01.
[54] X. Pan, Y. Cao, X. Du, B. He, G. Fang, R. Shao, and Y. Chen, "Flowcog: Context-aware semantics extraction and analysis of information flow leaks in Android apps," in *Proc. 27th USENIX Security Symposium, (SEC'18)*, 2018.
[55] R. Pandita, X. Xiao, W. Yang, W. Enck, and T. Xie, "WHYPER: towards automating risk assessment of mobile applications," in *Proc. 22th USENIX Security Symposium (SEC'13)*, 2013.
[56] Prolific, "Prolific," https://www.prolific.com, accessed: 2024-07-01.
[57] Z. Qu, V. Rastogi, X. Zhang, Y. Chen, T. Zhu, and Z. Chen, "Autocog: Measuring the description-to-permission fidelity in Android applications," in *Proc. 21st ACM Conference on Computer and Communications Security (SIGSAC'14)*, 2014.
[58] D. Rodriguez, A. Jain, J. M. D. Alamo, and N. Sadeh, "Comparing privacy label disclosures of apps published in both the app store and google play stores," in *Proc. 8th IEEE European Symposium on Security and Privacy Workshops (EuroS&PW'23)*, 2023.
[59] G. L. Scoccia, M. Autili, G. Stilo, and P. Inverardi, "An empirical study of privacy labels on the apple ios mobile app store," in *Proc. 9th IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft '22)*, 2022.
[60] A. Sela, S. C. Wheeler, and G. Sarial-Abi, "We are not the same as you and i: Causal effects of minor language variations on consumers' attitudes toward brands," *Journal of Consumer Research*, vol. 39, pp. 644–661, 2012.
[61] B. Shen, L. Wei, C. Xiang, Y. Wu, M. Shen, Y. Zhou, and X. Jin, "Can systems explain permissions better? understanding users' misperceptions under smartphone runtime permission model," in *Proc. 30th USENIX Security Symposium, (SEC'21)*, 2021.
[62] P. F. Stalmeier, M. S. Roosmalen, L. C. Verhoef, J. E. Hoekstra-Weebers, J. C. Oosterwijk, U. Moog, N. Hoogerbrugge, and W. A. van Daal, "The decision evaluation scales," *Patient Education and Counseling*, vol. 57, pp. 286–293, 2005.
[63] A. Steiger and A. Kühberger, "A meta-analytic re-appraisal of the framing effect," *Zeitschrift für Psychologie*, vol. 226, pp. 45–55, 2018.
[64] A. Stopper and J. Caltrider, "See no evil: Loopholes in google's data safety labels keep companies in the clear and consumers in the dark," https://foundation.mozilla.org/en/campaigns/googles-data-safety-labels/, accessed: 2024-07-01.
[65] J. Tan, K. Nguyen, M. Theodorides, H. Negrón-Arroyo, C. Thompson, S. Egelman, and D. A. Wagner, "The effect of developer-specified explanations for permission requests on smartphone user behavior," in *Conference on Human Factors in Computing Systems (CHI'14)*, 2014.
[66] A. Tversky and D. Kahneman, "The framing of decisions and the psychology of choice," *Science*, vol. 211, pp. 453–458, 1981.
[67] H. Wang, J. I. Hong, and Y. Guo, "Using text mining to infer the purpose of permission use in mobile apps," in *Proc. ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'15)*, 2015.
[68] M. Ward and A. W. Meade, "Dealing with careless responding in survey data: Prevention, identification, and recommended best practices," *Annual Review of Psychology*, vol. 74, pp. 577–596, 2023.
[69] K. Watson, M. Just, and T. Berg, "A comic-based approach to permission request communication," *Comput. Secur.*, vol. 124, p. 102942, 2023.
[70] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. A. Wagner, and K. Beznosov, "Android permissions remystified: A field study on contextual integrity," in *Proc. 24th USENIX Security Symposium (SEC'15)*, 2015.
[71] Y. Xiao, Z. Li, Y. Qin, X. Bai, J. Guan, X. Liao, and L. Xing, "Lalaine: Measuring and characterizing non-compliance of apple privacy labels," in *Proc. 32nd USENIX Security Symposium (SEC'23)*, 2023.
[72] B. Zhang and H. Xu, "Privacy nudges for mobile applications: Effects on the creepiness emotion and privacy attitudes," in *Proc. 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW'16)*, 2016.
[73] M. Zhang, Y. Duan, Q. Feng, and H. Yin, "Towards automatic generation of security-centric descriptions for Android apps," in *Proc. 22nd ACM Conference on Computer and Communications Security (SIGSAC'15)*, 2015.
[74] S. Zhang, Y. Feng, Y. Yao, L. F. Cranor, and N. Sadeh, "How usable are ios app privacy labels?" *Proc. Priv. Enhancing Technol.*, vol. 4, pp. 204–228, 2022.
[75] S. Zhang and N. Sadeh, "Do privacy labels answer users' privacy questions?" *Workshop on Usable Security and Privacy*, 2023.

APPENDIX A
DEMOGRAPHICS OF PARTICIPANTS

Table III presents the demographic information of our 960 participants, while Table IV shows their countries of residence.

APPENDIX B
QUESTIONNAIRE

In this section, you will find the survey questions, keeping in mind that the questions in Section B-A are repeated for each permission (camera, location, storage, and microphone).

A. Rationale Questions

On the next pages, you'll find four messages from various smartphone apps. Pretend you've just installed these apps on your phone and the first thing you see upon opening each app is one of these messages. We'd like you to react to the message and then answer a few questions to share your opinion.

TABLE III: Participants' demographics

Number of Participants	**960**	
Gender		
Male	478	49.8%
Female	466	48.5%
Other	14	1.5%
Prefer not to say	2	0.2%
Age		
18–24	269	28.0%
25–34	416	43.3%
35–44	151	15.7%
45–54	77	8.0%
55–64	36	3.8%
65 and over	11	1.1%
Education		
Less than high school	4	0.4%
High school	133	13.9%
Some college	164	17.1%
Associate degree	61	6.4%
Bachelor degree	421	43.9%
Master degree	141	14.7%
Doctoral degree	14	1.5%
Professional degree	16	1.7%
Other	6	0.6%
Mobile OS		
Android	579	60.3%
iOS	366	38.1%
Windows	15	1.6%

TABLE IV: Participants' country of residence

Country of Residence	**960**	
Africa	**119**	**12.4%**
South Africa	119	12.4%
Americas	**358**	**37.3%**
United States	131	13.6%
Mexico	115	12%
Canada	98	10.2%
Chile	14	1.5%
Asia	**10**	**1.0%**
Israel	7	0.7%
Japan	2	0.2%
Korea	1	0.1%
Europe	**433**	**45.1%**
Poland	100	10.4%
Portugal	94	9.8%
Italy	56	5.8%
Spain	37	3.9%
United Kingdom	32	3.3%
Greece	26	2.7%
Hungary	21	2.2%
Estonia	11	1.1%
Latvia	8	0.8%
Netherlands	8	0.8%
Czech Republic	7	0.7%
Finland	6	0.6%
France	5	0.5%
Ireland	4	0.4%
Slovenia	4	0.4%
Switzerland	4	0.4%
Belgium	3	0.3%
Germany	2	0.2%
Sweden	2	0.2%
Austria	1	0.1%
Denmark	1	0.1%
Norway	1	0.1%
Oceania	**40**	**4.2%**
Australia	40	4.2%

Imagine that you just installed the following app on your phone. Please continue to see the next screen of this app.

[Show Figure 13a]

After opening the app, it shows you the following message. Please take a moment to carefully read this message.

[Show Figure 13b]

Decision: Would you allow this app access to your {*permission*}? Please carefully read the message in the screenshot before making your decision.

- ○ Allow
- ○ Deny

Decision Satisfaction: In a previous question, you decided to {*allow/deny*} the app access to your {*permission*}. Please indicate your agreement with the following statements concerning your decision: [1–strongly disagree 2–disagree 3–somewhat disagree 4–neither agree nor disagree 5–somewhat agree 6–agree 7–strongly agree]

- I expect to stick with my decision.
- I am satisfied with my decision.
- I am doubtful about my choice.
- I would make the same decision if I had to interact with this app again.
- I am very confident that I made the right decision for myself.

Informed Decision: Please indicate your agreement with the following statements concerning your decision: [1–strongly disagree 2–disagree 3–somewhat disagree 4–neither agree nor disagree 5–somewhat agree 6–agree 7–strongly agree]

- I am satisfied with the information I received.
- I know the pros and cons of granting this app access to my {*permission*}.
- I would have liked more information about how the app will use the access to my {*permission*}.
- I made a well-informed choice.
- I know exactly why the app needs access to my {*permission*}.

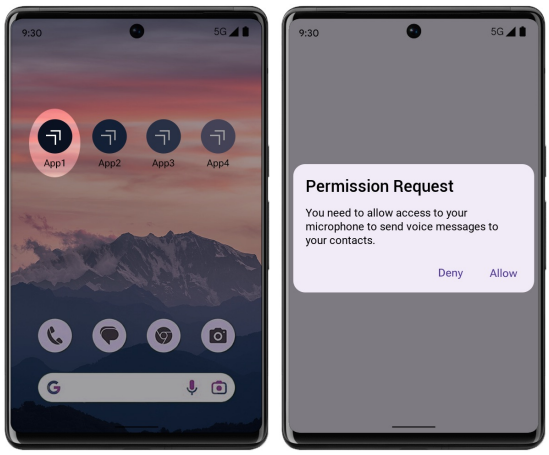

(a) Home screen (b) Rationale screen

Fig. 13: Sample screens from the questionnaire.

Decision Control: Please indicate your agreement with the following statements concerning your decision: [1–strongly disagree 2–disagree 3–somewhat disagree 4–neither agree nor disagree 5–somewhat agree 6–agree 7–strongly agree]

- I felt pressured by the app to make this decision.
- The app allowed me to make my own decision.
- I feel that the app forced me to make this decision.
- This was my own decision.
- I felt that the app would exclude me from using its features if I refused to grant access to my {*permission*}.

B. Demographic Questions

Gender: Which gender do you identify most with?

- ○ Male
- ○ Female
- ○ Prefer not to say
- ○ Other _____

Year of Birth: What is your year of birth? _____

Education: What is the highest level of school you have completed or the highest degree you have received?

- ○ Less than high school degree
- ○ High school graduate (high school diploma or equivalent including GED)
- ○ Some college but no degree
- ○ Associate degree in college (2–year)
- ○ Bachelor's degree in college (4–year)
- ○ Master's degree
- ○ Doctoral degree
- ○ Professional degree (JD, MD)
- ○ Something else, namely: _____

Mobile OS: What operating system are you using on your (primary) mobile phone?

- ○ Android
- ○ iOS (iPhone)
- ○ Windows (Windows Phone)
- ○ Other _____

Privacy Concerns: Please rate the following statements: [1–strongly disagree 2–disagree 3–somewhat disagree 4–neither agree nor disagree 5–somewhat agree 6–agree 7–strongly agree]

- Compared to others, I am more sensitive about the way mobile apps handle my personal information.
- To me, it is the most important thing to keep my privacy intact from mobile apps.
- In general, I am very concerned about threats to my personal privacy.

Prior Privacy Experience: Please answer the following questions: [1–never 2–very rarely 3–rarely 4–occasionally 5–frequently 6–very frequently 7–always]

- How often have you personally experienced incidents whereby your personal information was used by some mobile app without your authorization?
- How much have you heard or read during the last year about the use and potential misuse of the information collected from mobile apps?
- How often have you personally been the victim of what you felt was an improper invasion of your privacy from a mobile app?

Use My Answers: Sometimes, when people take part in a survey, they may not pay full attention or get distracted. When this happens, the answers they provide may not be good for scientific research. Please answer honestly: Were you diligent and attentive when answering this questionnaire? You will be paid for your participation regardless of your answer.

- ○ No
- ○ Yes

APPENDIX C
DES ITEM FIT AND CONSISTENCY

We calculated multilevel internal consistency using McDonald's Omega [27] and determined the factor loadings for all items in the Decision Evaluation Scales (DES). Table V presents the results for all items and subscales of the DES. The newly included Informed Decision and Decision Satisfaction items showed a very good fit. In contrast, the item added to Decision Control had a somewhat weaker fit but was still considered acceptable. Despite variations in factor loadings, all subscales demonstrated good internal consistency within and across participants, indicating that the scales' consistency remained robust.

APPENDIX D
MODEL FIT

We statistically compared all steps in the model-building process using the akaike information criterion (AIC) and the likelihood-ratio tests. The model that described our data best and had the lowest AIC score was selected as the final model. Table VI presents the goodness of fit, marginal R^2, and conditional R^2 for each step in the model-building process of all outcome variables.

TABLE V: Standardized item fit and internal consistency measures for the DES subscales.

	DES Inform	DES Satis	DES Control
I am satisfied with the information I received.	0.87		
I know the pros and cons of granting this app access to my {permission}.	0.58		
I would have liked more information about how the app will use the access to my {permission}.	0.60		
I made a well-informed choice.	0.69		
I know exactly why the app needs access to my {permission}.	0.82		
I expect to stick with my decision.		0.81	
I am satisfied with my decision.		0.88	
I am doubtful about my choice.		0.72	
I would make the same decision if I had to interact with this app again.		0.80	
I am very confident that I made the right decision for myself.		0.88	
I felt pressured by the app to make this decision.			0.85
The app allowed me to make my own decision.			0.49
I feel that the app forced me to make this decision.			0.88
This was my own decision.			0.45
I felt that the app would exclude me from using its features if I refused to grant access to my {permission}.			0.45
$\omega_{between}$	0.84	0.93	0.84
ω_{within}	0.87	0.89	0.77

TABLE VI: Goodness of fit for final models

Decision	AIC	LogLik	Df	Pr(>Chisq)	Marginal R^2	Conditional R^2
simple regression	5050.1	-2524.0				
step 1: multilevel base (user and permission as random effects)	4963.5	-2478.8	2	<0.001		0.161
+ step 2: variables from previous work	4882.5	-2436.3	2	<0.001	0.041	0.167
+ step 3: variables of interest: rationale building blocks	4874.3	-2423.1	9	0.002	0.050	0.177
+ step 4: interactions between building blocks	4905.6	-2411.8	27	0.702	0.058	0.188
DES Inform						
simple regression	13661.7	-6828.8				
step 1: multilevel base (user and permission as random effects)	13252.7	-6622.4	2	<0.001		0.299
+ step 2: variables from previous work & Decision	12642.9	-6314.4	3	<0.001	0.133	0.414
+ step 3: variables of interest: rationale building blocks	12621.0	-6294.5	9	<0.001	0.142	0.413
+ step 4: interactions between building blocks	12667.4	-6290.7	27	0.999	0.143	0.414
DES Satis						
simple regression	11489.1	-5742.5				
step 1: multilevel base (user and permission as random effects)	10903.8	-5447.9	2	<0.001		0.364
+ step 2: variables from previous work & Decision	10800.1	-5393.0	3	<0.001	0.033	0.377
+ step 3: variables of interest: rationale building blocks	10792.5	-5380.2	9	0.002	0.037	0.381
+ step 4: interactions between building blocks	10828.0	-5371.0	27	0.889	0.041	0.384
DES Control						
simple regression	12128.8	-6062.4				
step 1: multilevel base (user and permission as random effects)	11036.5	-5514.2	2	<0.001		0.496
+ step 2: variables from previous work & Decision	10990.7	-5488.4	3	<0.001	0.025	0.496
+ step 3: variables of interest: rationale building blocks	10982.7	-5475.4	9	0.002	0.029	0.499
+ step 4: interactions between building blocks	11017.0	-5465.5	27	0.844	0.032	0.503