

Beyond the Accept Button: How Information and Control Shape Data Sharing and AI Engagement

IBRAHIM CHHAYA*, Carnegie Mellon University, USA

YASH MAURYA*, Carnegie Mellon University, USA

ZUOFEI HONG*, Carnegie Mellon University, USA

LIMIN GE*, Carnegie Mellon University, USA

This study examines how consent flow design impacts user engagement with AI features and data sharing in digital platforms. Through a simulated social media app onboarding experience as a part of our survey with 500 people, we investigated the effects of content length, information presentation, and granularity of data sharing controls. Our findings reveal three key insights: (1) users found longer and medium-length content more comprehensible than shorter variants, suggesting detailed information improves clarity when properly presented; (2) although provided consent flow information did not impact user sentiment or behavior, participants strongly desired more transparency about third-party data sharing and AI training practices; and (3) interface design significantly influenced AI privacy choices, with longer interfaces promoting basic privacy settings while compact interfaces encouraged more granular control.

1 Introduction

Digital platforms use consent flows to help users make decisions about data sharing, privacy, and engagement with platform features. These interfaces need to balance multiple competing needs: providing clear information, ensuring user comprehension, and maintaining engagement while complying with privacy regulations such as the European Union General Data Protection Regulation's requirement for consent to be for "one or more specific purposes" [15]. Even when not legally mandated, privacy best practices suggest offering users granular privacy controls on how specific types of information are used and shared [2]. However, designing effective consent flows remains challenging, particularly as platforms introduce more complex features such as data sharing between applications and AI-powered capabilities.

Research suggests that users often experience difficulty engaging meaningfully with consent interfaces [9, 12]. When presented with dense or complex privacy information, users may make

*Equal contribution to this research.

Authors' Contact Information: Ibrahim Chhaya, Carnegie Mellon University, Pittsburgh, PA, USA, ichhaya@andrew.cmu.edu; Yash Maurya, Carnegie Mellon University, Pittsburgh, PA, USA, ymaurya@andrew.cmu.edu; Zuofei Hong, zuofeih@andrew.cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA; Limin Ge, liminge@andrew.cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA.

decisions that do not align with their actual preferences or data sharing goals. This misalignment can be due to various factors, including information presentation, complexity of the interface, and the structure of privacy controls.

Recent work has examined how interface design elements, such as the number of screens, settings per screen, and bundling approaches affect user engagement with privacy choices. Following Habib and Cranor’s privacy choice evaluation framework [5], Alashwali et al. [1] found that while users preferred privacy choices split across multiple screens, the actual number of screens and bundling approaches had limited impact on user behavior. Their study of a social media app’s onboarding experience highlighted how certain interface elements, particularly preset options, could significantly influence privacy decisions.

Building on this work, our study examines how consent flow design affects user engagement, which is defined as user interaction with features, through a simulated cross-app scenario involving BuyNow (shown in Figure 1), a fictitious marketplace app, and its social media extension RightNow (shown in Figure 2). Through a survey of 500 participants, we investigated three key aspects: how content length and organization affect comprehension, sentiment and behavior; what information users find valuable in consent flows; and how different levels of control impact willingness to engage with AI features.

Our investigation revealed several findings about consent flow design. We found that the length and *presentation of information* influenced user comprehension and behavior. Although users consistently expressed a desire for more transparency about third-party data sharing and AI training practices, the *information provided* had minimal impact on their behavior. The interface design also showed notable effects on how users approached *AI data-sharing privacy controls*.

These findings provide empirical insights into how consent flow design elements affect user behavior and their interaction AI data-sharing privacy controls. Our results may help inform the development of more effective consent flows that balance information presentation with user control, particularly for platforms implementing AI features and cross-platform data sharing.

1.1 Research Objectives

Building on prior work in consent interface design, this study investigates how specific design elements affect user engagement with privacy choices in cross-app and AI-enabled contexts. We aim to answer the following research questions:

- **RQ1:** How do the number of pages and explanatory information in the consent flow impact user sentiment, comprehension, and behavior?
- **RQ2:** What information would users like to have as part of the consent flow, and how does having or not having this information affect their sentiment, comprehension, and behavior?

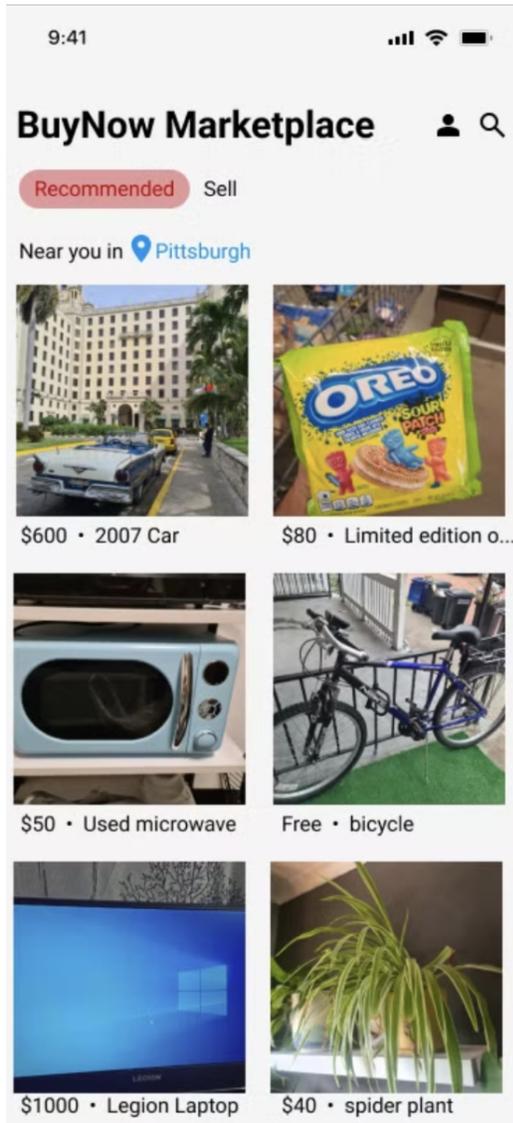


Fig. 1. BuyNow App Interface

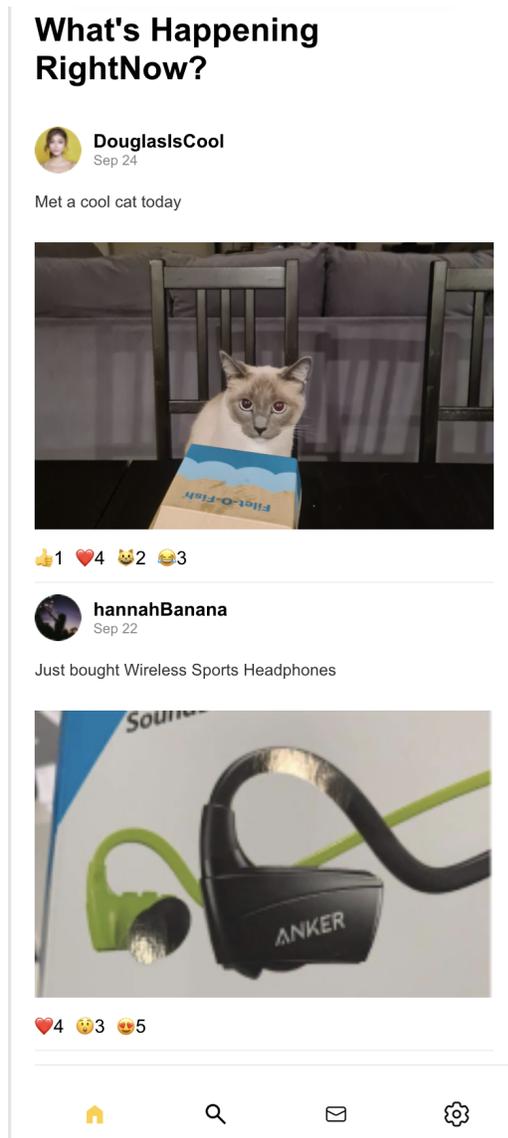


Fig. 2. RightNow App Interface

- **RQ3:** How do different levels of control over data sharing affect users' willingness to engage with AI features in a social media platform?

2 Related Work

2.1 Interface Design for Privacy Controls

The complexity of modern privacy controls is evident on the main social media platforms. For example, X (formerly Twitter) offers eleven menu options within its “Privacy and safety” settings, while Facebook provides eight categories of “Audience and visibility” settings and four categories of “Ad preferences” settings.

This proliferation of options presents significant usability challenges. Multiple studies have identified the difficulties users face in locating and utilizing privacy controls. Hsu et al. discovered that most of the participants were unaware of Facebook’s ad personalization settings [8]. This finding was further supported by the research by Habib et al., which revealed a widespread difficulty in finding settings “Ad Preferences” [7]. In a separate study, Habib et al. also found that most participants required help finding opt-out settings, with many users expressing confusion over poorly labeled options [6]. Furthermore, several studies on social media privacy settings have documented mismatches between users’ actual privacy settings and their sharing intentions, often resulting in unintended information disclosure [10, 11, 13].

Research by Alashwali et al. complements these findings by exploring how different interface design elements, such as the number of screens and the grouping of settings, affect users’ ability to navigate privacy settings. Their study found that splitting settings across multiple screens and grouping them functionally (e.g., notifications, ads, or privacy invasiveness) can improve usability and reduce fatigue [1]. This suggests that structuring privacy controls to match user expectations can mitigate cognitive overload and improve transparency.

These studies reveal common problems users face in privacy control interfaces and provide a direction for improving design solutions for our research by confirming users’ strong demand for privacy transparency and data control. These insights support our research objective to investigate the impact of interface content length and granularity on users’ willingness to share data and engage with AI functions.

2.2 User Information Requirements and Privacy Decision-Making

The complexity of privacy controls has significant implications for user behavior and decision-making. Keith et al. found that complex privacy controls can increase the disclosure of personal data, attributing this phenomenon to privacy fatigue [9]. This finding was corroborated by Choi et al., who demonstrated that privacy fatigue strongly influences increased disclosure of personal data and user disengagement [3]. A comprehensive review by Schyff et al. examined previous research on privacy fatigue, linking it primarily to users’ perceived loss of control over their personal data [14]. Alashwali et al. found that inline presentation of detailed information for each

setting (as opposed to requiring users to click for more details) can enhance understanding and support informed decision-making [1]. Their study also revealed that participants preferred to make privacy decisions upfront during the app onboarding process, emphasizing the value of timely and relevant information.

Another critical finding from Alashwali et al.'s study was the unintended effects of preset options in privacy interfaces. While presets can reduce user effort, they often nudge users toward less privacy-protective choices and lead to reduced comprehension of their selections [1].

These papers inform our research, highlighting the need to evaluate how different levels of information detail and control granularity affect users' trust, understanding, and engagement with privacy settings. Privacy interfaces should address users' concerns about transparency and control while minimizing information overload.

2.3 Evaluation of Privacy Study Designs

Prior research in usable privacy by Habib and Cranor [5] has established several approaches for evaluating privacy interface designs. These evaluations cover three metrics: sentiment (measuring user fatigue, and perceived clarity), comprehension (how well users understand privacy policies and features), and behavioral metrics (analyzing user interactions and choices). Studies have particularly examined how notice and consent flow design elements like content length, information presentation, and granularity of controls affect user engagement.

Distler et al. [4] show that evaluating privacy interfaces presents unique challenges, particularly in balancing ecological validity with practical constraints. Lab studies can help control external variables but may not fully capture real-world privacy behaviors and perceptions. When participants are given fictitious information for privacy studies, they may behave differently than they would with their own personal data at risk. Additionally, the frequency of simulated privacy events in lab settings often differs from real-world occurrence rates, which can affect ecological validity.

Our research builds on these evaluation approaches by combining controlled experiments with survey-based measurements. Through our simulated social media app onboarding experience, we systematically evaluated how different aspects of consent flow design impact user engagement, comprehension, and privacy decisions.

3 Methods

We conducted a 500-participant online survey to compare nine variants (shown in Table 1) of the onboarding interfaces of a fictional social media app and explore the impact of the number of pages and explanatory information in the consent flow, and what information users would like as part of the consent flow on user sentiment, comprehension, and behavior, as well as how different

levels of control affect user willingness to engage with AI features. It should be noted that the wording presented on the screens was not the focus of our research.

Table 1. Variants: Content Length and AI Control Variations

Length	Group Name	Control of AI
Short one page	Short 1 (S1)	Training AI notice
	Short 2 (S2)	Training AI opt in/opt out
	Short 3 (S3)	Training AI partial opt out
Medium text	Medium 1 (M1)	Training AI notice
	Medium 2 (M2)	Training AI opt in/opt out
	Medium 3 (M3)	Training AI partial opt out
Long pages	Long 1 (L1)	Training AI notice
	Long 2 (L2)	Training AI opt in/opt out
	Long 3 (L3)	Training AI partial opt out

3.1 Study Design

The study was conducted using a simulated onboarding process involving the fictional marketplace app BuyNow and the fictional social media app RightNow, which resemble real-life marketplace apps such as Facebook Marketplace, Etsy and OfferUp, and real-life social media apps such as Facebook, Instagram and Twitter / X respectively. BuyNow is defined as a marketplace app whose data could be used to enhance the experience on a new social media platform. RightNow is defined as a social media app launched by the same company, which allows users to seamlessly integrate with their BuyNow account details.

We designed our RightNow user interfaces in Figma and implemented them in the React web library, designed to resemble a mobile app layout. The application logged user interactions on our server, including if users clicked to receive more information, their choice selections, the page navigated, and the time spent per screen.

The RightNow application is a continuation of the work of Alashwali et al. [1] on consent interface design. Alashwali et al. [1] focused on designing and evaluating onboarding interfaces for a single app, exploring how variations in the number of screens, bundling strategies, and preset options affected user sentiment, comprehension, and privacy decision making. Their study highlighted user preferences for multi-screen designs and identified the influence of preset options on behavior.

In contrast, our study examines cross-app scenarios involving BuyNow and RightNow, reflecting real-world systems where data is shared across interconnected platforms. By implementing nine onboarding variants that varied in content length and control granularity, we investigated the

Beyond the Accept Button: How Information and Control Shape Data Sharing and AI Engagement

effects of these design elements on user engagement, comprehension, and sentiment in a multi-app ecosystem.

3.2 RightNow Interface Variants

The RightNow application implements a controlled experimental design with multiple variants to study social media onboarding experiences, featuring both constant and variable interface elements. At its core, the application maintains three standardized screens in all variants, including a contact synchronization request interface, as shown in Figure 3. Depending on the experimental condition, participants encounter three or four distinct settings related to social media onboarding in total, with the additional screens varying in their interface design to address specific research questions and mirror real-world social media application patterns.

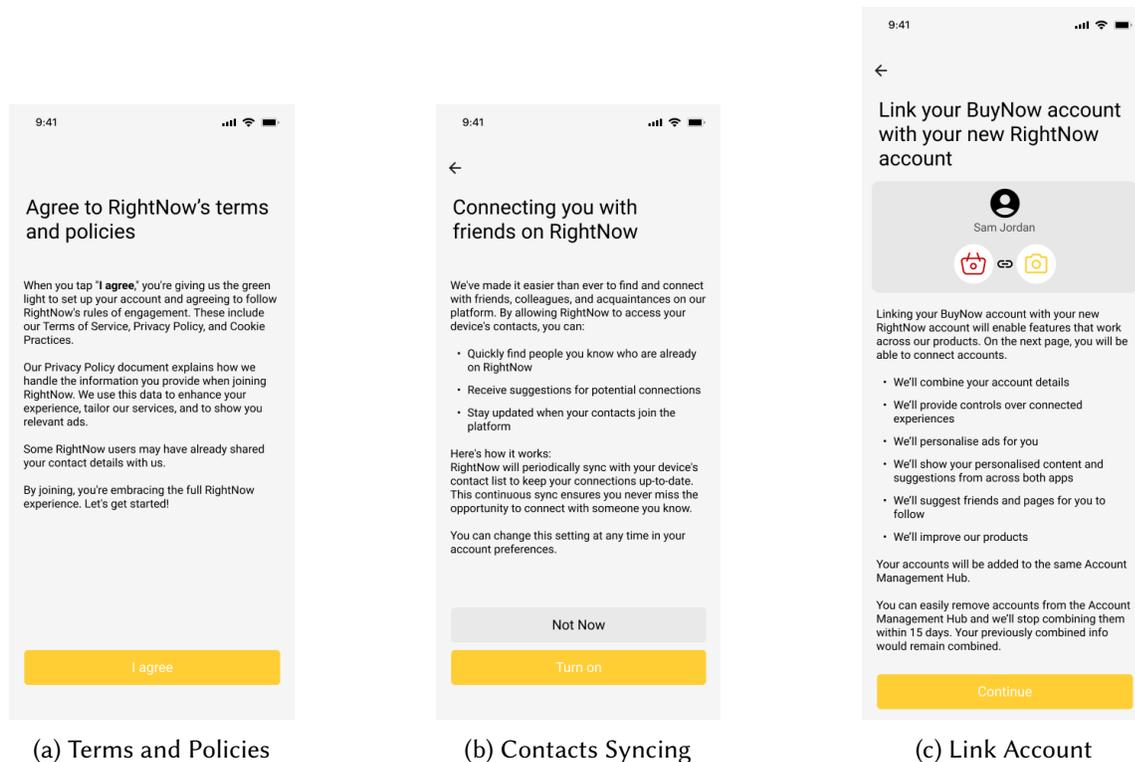


Fig. 3. Standard interfaces shown across all variants

The RightNow interface employs a variation in both information presentation and user control through nine distinct variants, organized in a 3×3 matrix of length and control options. Initial information visibility varies across the three length categories (*Short*, *Medium*, and *Long*) with content hidden behind an *expandable accordion* mechanism, shown in Figure 4. Although all

variants contain identical information, the Short and Medium variants require user interaction to access additional content through the accordion interface, which expands upon user tap. The Long variant displays all information upfront, distinguishing it from its counterparts by eliminating the need for user-initiated content expansion, as can be seen in Figure 5.

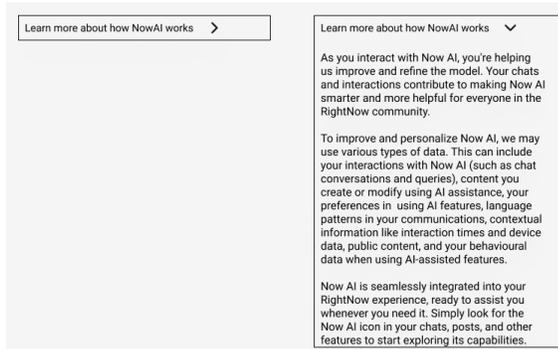


Fig. 4. Expandable accordion showing closed state on the left and opened on the right

The layout of the pages also differed for each variant: Short had all settings on one page with no extra information displayed outside the expandable accordion, Medium presented each setting on its own page with some information displayed outside the expandable accordion, while Long presented each setting on its own page with all of the information printed and no expandable accordion.

The subvariants differed in amount of control, particularly with the final setting: use of user data for AI training. Subvariant 1 (AI Training Notice) did not allow users to opt out of AI training, subvariant 2 (AI Opt In/Opt Out) presented users with an opt-in/opt-out selection, and subvariant 3 (AI Granular Opt Out) allowed users to select more granular settings: “Don’t use my data for training and only get limited NowAI model,” “Use my public data for training and get core NowAI features,” or “Use all interactions for training and get fully personalized NowAI model.” Screenshots of key differences between each variant are illustrated in Figure 6.

The variants and the AI subvariants were designed as part of our between-subjects study to answer the proposed research questions and develop design recommendations for the project sponsor Meta. The three length variants (Short, Medium, and Long) are layouts typically seen in real-life app onboarding experiences and enable us to test length and amount of information per screen. The AI subvariants were designed to test level of control and willingness to engage with AI features, which is especially important to the sponsor as both AI technologies and privacy awareness continue to become commonplace.

The experiment was designed with the following independent variables:

Beyond the Accept Button: How Information and Control Shape Data Sharing and AI Engagement

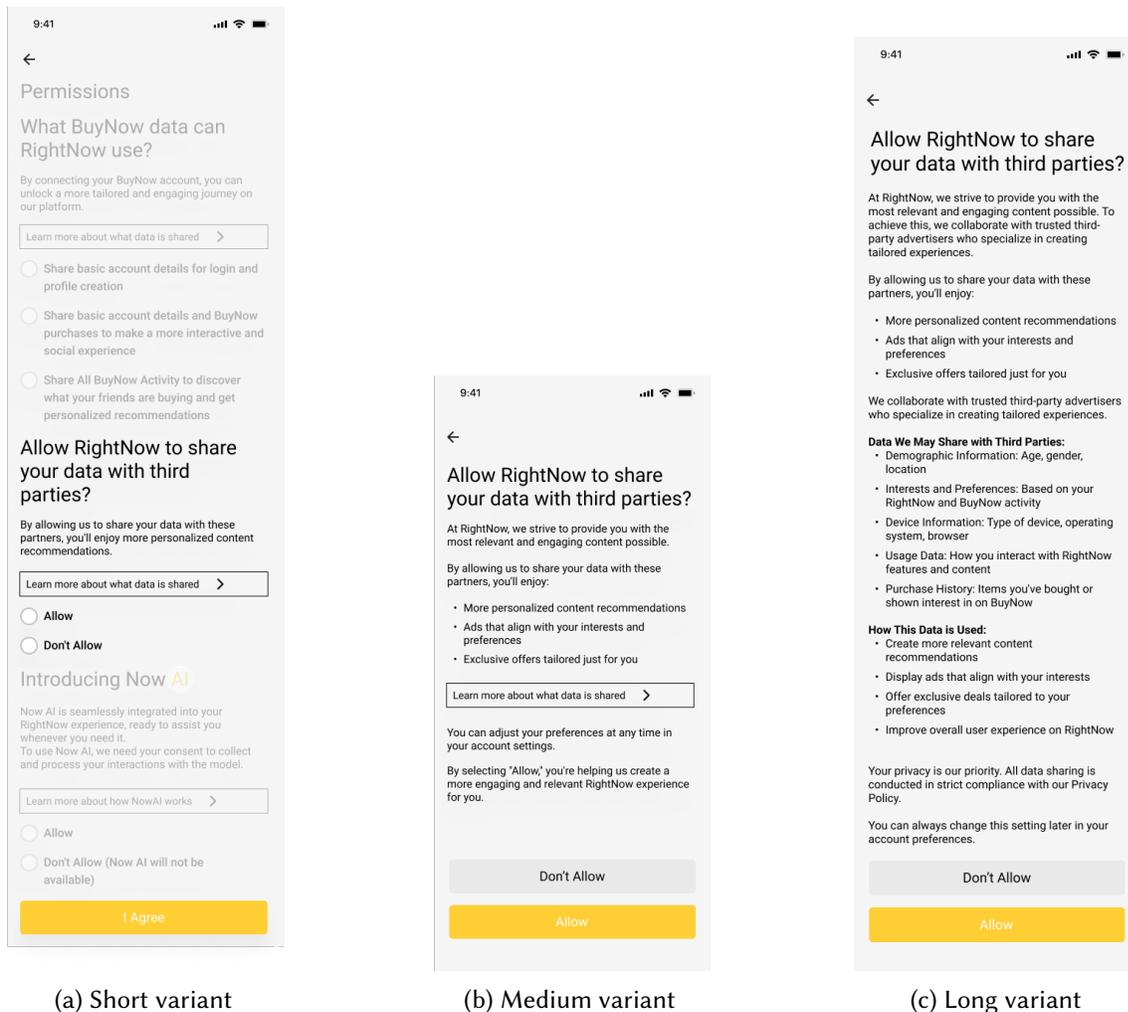


Fig. 5. Variant interfaces showing difference in amount of information visible at first. Figure 5a has a minimal amount of information visible, Figure 5b has a balanced amount of information visible, and Figure 5c has all the information visible. All these variants have the exact same information when the accordions are expanded.

- **Consent Flow Length:** Three variations (shown in Figure 5) — Short, Medium, and Long.
- **Control Over Data Sharing:** Three levels of control (shown in Figure 6): Notice, Opt-in/Opt-out, and Granular Opt-out.

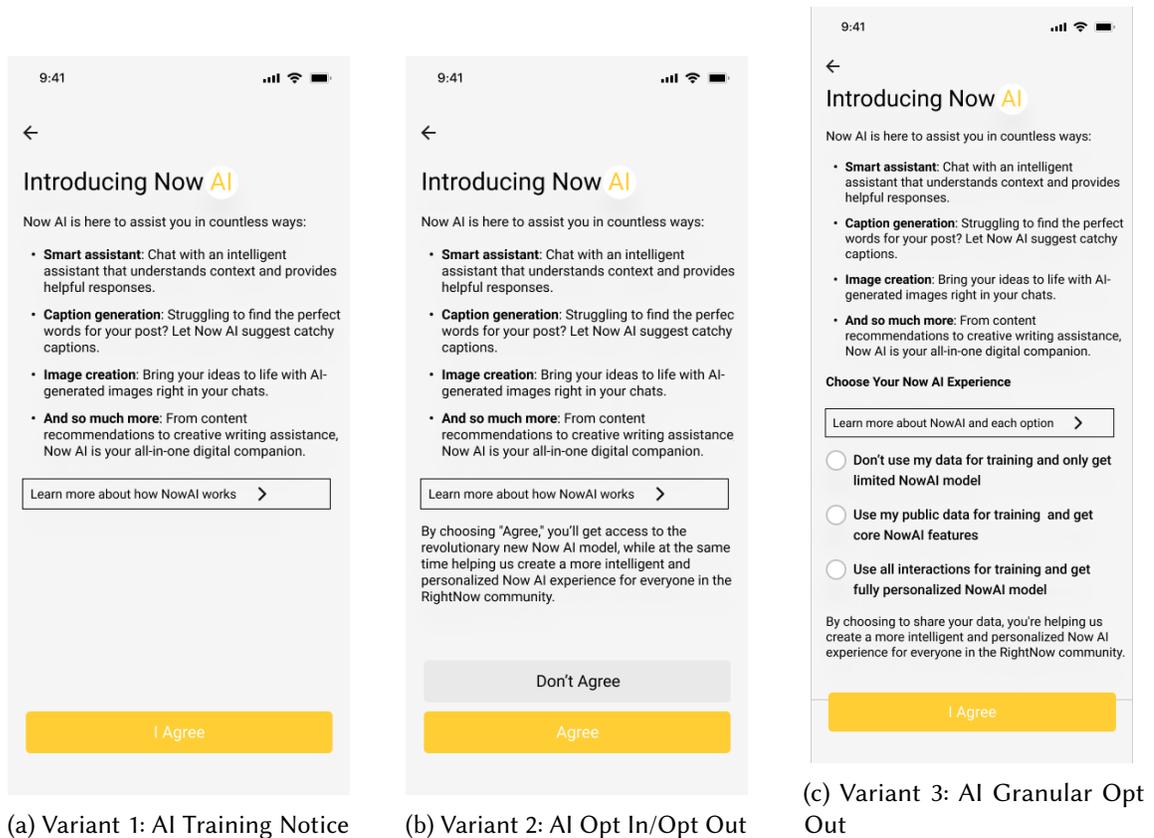


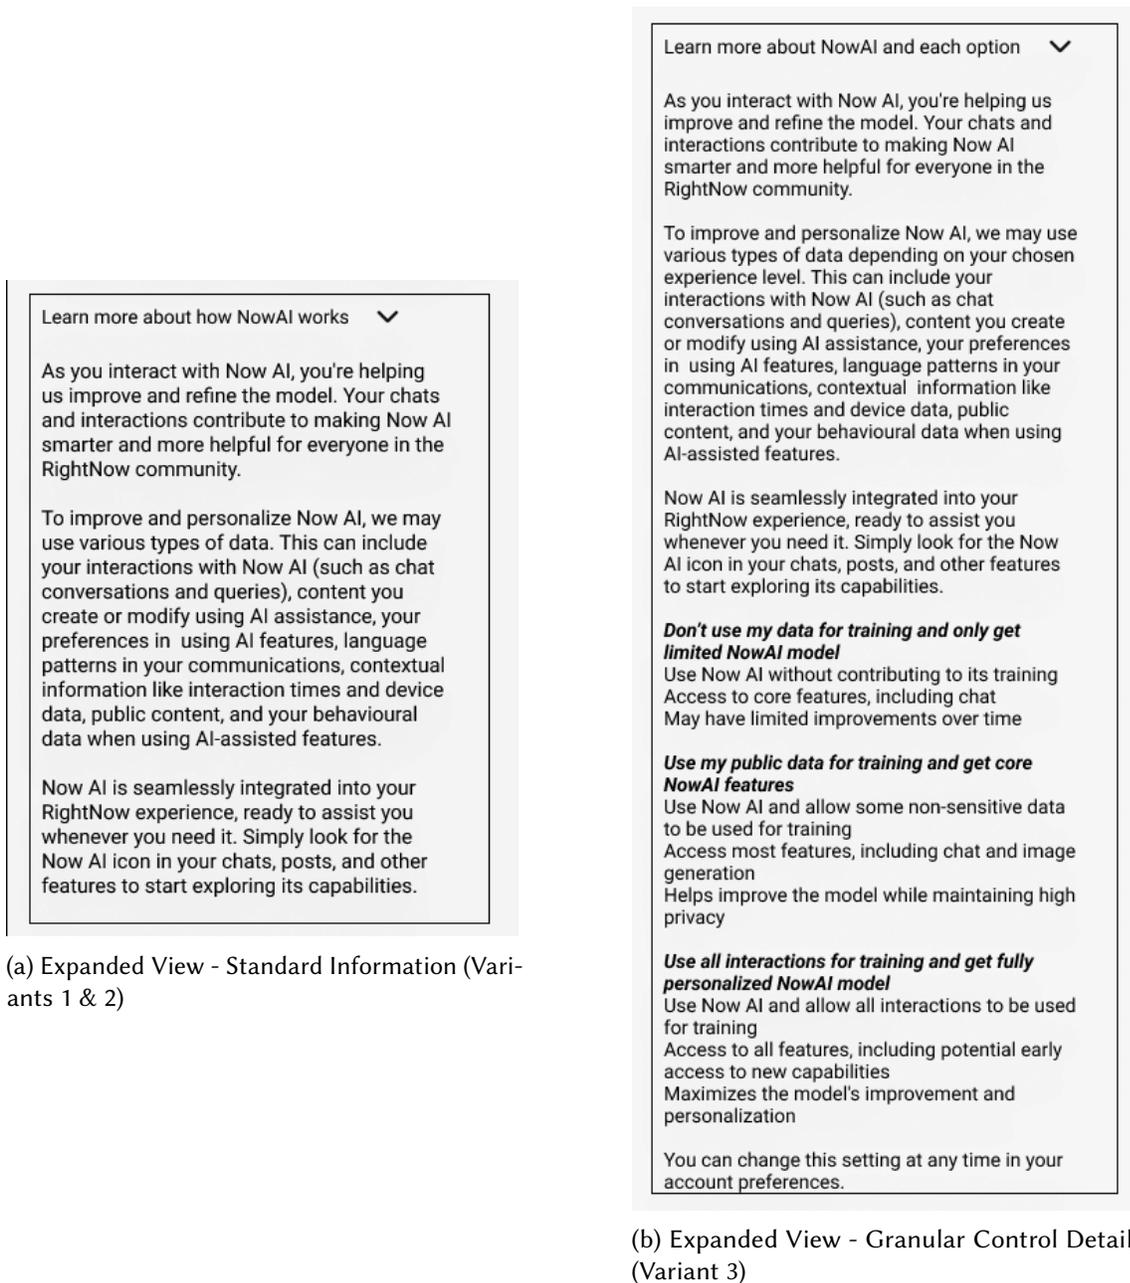
Fig. 6. Initial states of three AI training consent interface variants. Variant 1 (left) provides only information, Variant 2 (middle) offers a simple binary choice, and Variant 3 (right) presents granular controls for specific AI features. These variants demonstrate increasing levels of user control over AI training data usage. Expanded accordions for each screen shown in Figure 7

3.3 Survey Design

The survey was implemented in Qualtrics, containing multiple-choice and open-ended questions. The survey began with a consent form and scenario information, followed by a link to our simulated social media app onboarding experience. Participants were instructed to select the settings they would have chosen if they had signed up for this app in real life to observe behavior and collect real-life data. Despite this, we did not ask for or collect any personal information from the participants during the simulated app experience. Participants were randomly grouped into one of the 9 variants, while ensuring that each variant received a balanced distribution of participants.

After completing the simulated onboarding experience, participants were returned to the survey, which began with questions about participant social media and marketplace app usage to better

Beyond the Accept Button: How Information and Control Shape Data Sharing and AI Engagement



(a) Expanded View - Standard Information (Variants 1 & 2)

(b) Expanded View - Granular Control Details (Variant 3)

Fig. 7. Expanded states of the consent interface variants showing detailed information revealed by the accordions. Variants 1 and 2 share identical detailed information (left) focusing on general AI training practices, while Variant 3 (right) provides additional granular controls for specific AI features.

understand how active they are on apps that relate to our scenario. The participants were then asked questions about their sentiments, particularly about the amount of information presented, the level of control over data given, and their overall feelings. These were performed through a set of standard Likert-style questions on the perception of ease, length, fatigue, and lastly, level of control and amount of information. To examine comprehension and attention, we asked them a set of multiple-choice questions about the settings and features available in the app, such as "When can you change the setting related to contact syncing?" and "How can you manage the accounts if you want to unlink them in the future?" We also asked participants which options they selected during the onboarding experience (and why), whether the screens contained sufficient information, and if more controls were wanted.

For each participant, we collected their selected options and accordion clicks for each screen from the RightNow app logs which were connected by their prolific ids to the the survey responses. Figure 8 visualizes the data flow from participant recruitment through prolific to the results we present in our Results section after data analysis.

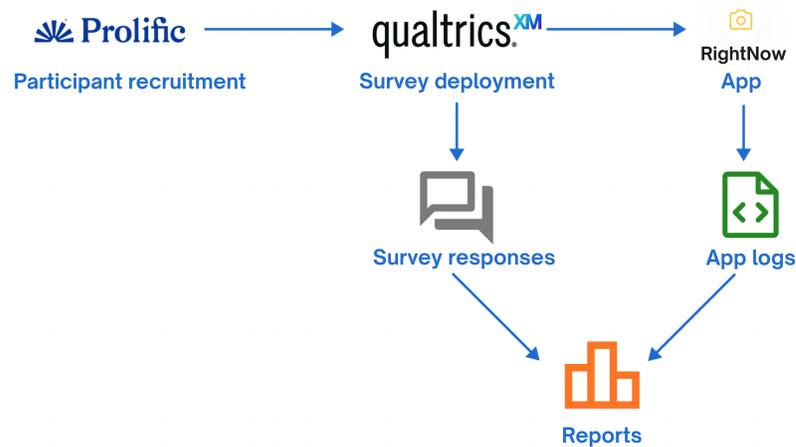


Fig. 8. Visualization of data flow from participant recruitment through survey completion and data processing

In addition, we asked them to select any settings that made them concerned about their privacy or about which they wanted more information, as well as how they typically select privacy-related settings for social media apps. In the end, we concluded the survey with a set of demographic questions. The complete survey is provided in Appendix D. S

3.4 Metrics and Data Collection

The study used the following metrics, aligned with the research questions:

- **Sentiment Metrics:** Self-reported fatigue, Satisfaction, and Perceived clarity (from survey responses).
- **Comprehension Metrics:** Correct identification of data-sharing policies and AI features (from survey responses).
- **Behavioral Metrics:** Click-through rates and opt-in / out options (from app logs).

Data was collected from two primary sources:

- **Survey Results:** The participants answered questions about their experience.
- **App Click Logs:** Interaction data from the simulated RightNow prototype.

Research Question	Metric	Data Source
RQ1	Fatigue, satisfaction, perceived clarity	Survey results
RQ2	Comprehension of settings and features	Survey results
RQ3	Opt-in/Opt-out behavior	Survey responses and App click logs

Table 2. Metrics aligned with research questions.

The study was conducted in the following phases:

- (1) **Pilot Study:** A pilot study with 23 participants was conducted to refine the survey questions and the consent flow design.
- (2) **Main Study:** The main study was conducted in three batches (20, 80, and 400 participants), each testing the experimental conditions.

3.5 Data Analysis

The data for our results and analysis were sourced from both the survey and the app interaction logs. The results were analyzed using descriptive statistics. We compared three sets of screens that differed in the amount of information visible at first (as explained in Section 3.2 and shown in Figure 5), as well as three sets of screens that differed in the granularity of the control (as explained in Section 3.2 and shown in Figure 6). To determine whether the distribution of participants between the variant groups was normal, we performed the Shapiro-Wilk test. We then analysed the results of the Likert-style questions using the Kruskal-Wallis rank sum test. If the results were statistically significant, we then performed pairwise tests using pairwise Wilcoxon rank sum tests with the p-values adjusted by the Bonferroni correction.

The categorical-style questions were analyzed using the Chi-squared test, and if significance was found, pairwise tests were performed using Fisher’s exact test and adjusted p-values using the Bonferroni correction. The same technique was used to test whether there were dependencies between the distribution of choices made by the participants in each variant.

In analyzing the open-ended responses received from the survey, two researchers collaborated to create a codebook for each question to categorize the responses into emergent themes. The researchers met in multiple sessions to resolve any ambiguous responses and decided on the final code for such responses by consensus.

3.6 Participant Recruitment

We first recruited 13 in-person participants and 10 Prolific participants for a pilot study to refine survey questions and the implementation of the application. For the main study, we recruited 500 participants for the 12-minute survey on social media onboarding experiences using Prolific, which included 2-4 minutes for the app onboarding. Our study screened participants using Prolific's screening tool who were at least 18 years old, residing in the United States of America, 50/50 male / female, fluent in English, and able to complete the study on a computer. The 500 participants were randomly assigned to each of the 9 variant groups, resulting in an actual average of 45 participants per variant. We published the study in three batches (20, 80, and 400 participants) to ensure error-free deployment.

Participants were compensated \$3.00 for completing the study, which is designed to be \$15.00 per hour on average.

3.7 Lessons Learned from Pilot Studies

Our pilot study revealed several opportunities to improve the study methodology. To improve data quality and participant experience, we simplified the survey by reducing free response fields and removing exploratory questions that were not central to our research questions. The consent form was refined with more generalized wording to avoid inadvertently influencing privacy-related responses. To improve comprehension of the simulated scenario, we enhanced visual clarity through bold text, larger fonts, and improved formatting. Finally, to better assess genuine understanding rather than familiarity with the interface, we implemented memory-based comprehension questions and discouraged participants from revisiting the app during the survey portion.

4 Results

We first cleaned the dataset by filtering out responses from participants who did not interact with the application (RightNow). This was tracked through app logs which included a participant's Prolific ID in the URL parameters. Participants were explicitly asked to complete the simulated onboarding experience before continuing with the survey and, as such, those who did not interact with the app were excluded from the study. Additionally, we filtered out participants who repeatedly interacted with the app. This was necessary because participants were instructed not to return

to the app after completing onboarding, as doing so could bias their responses in comparison to other participants.

We then analyzed the distribution of participants across the nine experimental conditions. The conditions varied by text length (Long, Medium, Short) and AI control type (Notice Only, General Opt-in/out, Granular Choices). The final sample consisted of 434 valid responses distributed as follows: Long conditions had 50 (L1), 45 (L2), and 58 (L3) participants; Medium conditions had 46 (M1), 42 (M2), and 53 (M3) participants; and Short conditions had 44 (S1), 45 (S2), and 51 (S3) participants.

Demographic data and social media use are provided in Tables 6, 7, and 8, and the distribution of the participants by variant is shown in Table 9. All of these tables are in Appendix B.

4.1 RQ1: How do the number of pages and explanatory information in the consent flow impact user sentiment, comprehension, and behavior?

4.1.1 *Impact on Sentiment.* To assess user sentiment across different consent flow designs (short, medium, and long), we analyzed responses to specific survey questions. Overall we found that participant sentiment towards all of our variants was fairly positive. The only statistically significant difference in answers to sentiment questions between variations was for the question: “How clear was the information provided on the screens regarding data sharing and AI features?” Participants found the medium variant more clear than the short variant.

Firstly, participants were asked about their fatigue during the sign-up process (Q6). Across all variants, responses were relatively similar: most participants were “Not fatigued at all” (S: 58.57%, M: 58.16%, L: 51.63%) or “Slightly fatigued” (S: 26.43%, M: 27.66%, L: 28.76%). This was followed by those who were “Moderately fatigued” (S: 11.43%, M: 10.64%, L: 12.42%), “Quite fatigued” (S: 3.57%, M: 2.84%, L: 4.58%), and “Very fatigued” (S: 0.00%, M: 0.71%, L: 2.61%). None of the distributions were normal and no statistically significant differences were observed ($p = 0.283$).

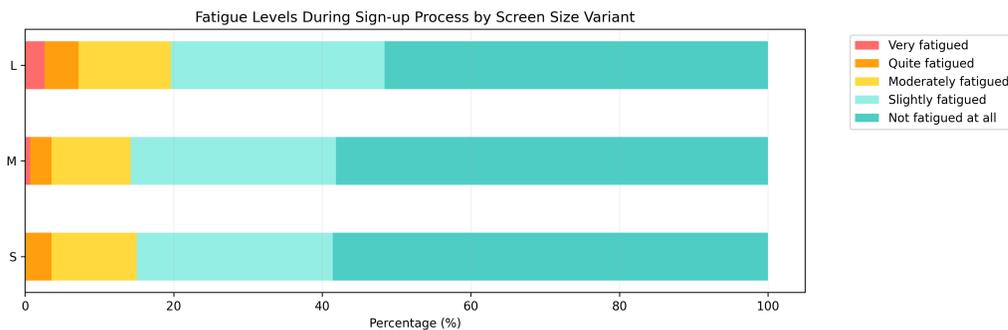


Fig. 9. Q6: Fatigue Felt By Variants

Secondly, participants were asked how likely they would be to complete the process if signing up for the RightNow app in real-life conditions (Q7). The majority indicated they would complete the process, with 66.01% to 78.56% selecting "Definitely" or "Likely." Completion rates were highest for variant S (78.56%), followed by M (68.08%) and L (66.01%). Negative responses ("Definitely not" or "Likely not") were lowest in variant S (10.00%) compared to M (17.02%) and L (15.68%). A Kruskal-Wallis test revealed moderate significance, but not statistically significant ($H = 4.23$, $p = 0.054$).

Participants were also asked to rate the ease of the sign-up process (Q9). Across all groups, the majority found the process straightforward, with 46.41% to 58.57% rating it as "Extremely easy." A smaller proportion found the process "Somewhat difficult" (S: 2.86%, M: 7.09%, L: 5.23%) or "Extremely difficult" (S: 0.00%, M: 0.71%, L: 0.65%). A Kruskal-Wallis test revealed no significant difference ($H = 4.23$, $p = 0.1208$).

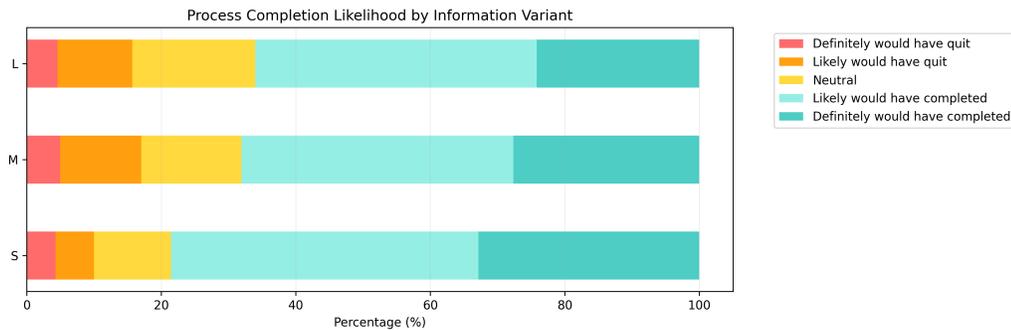


Fig. 10. Q7: App Sign-up Completion Likelihood

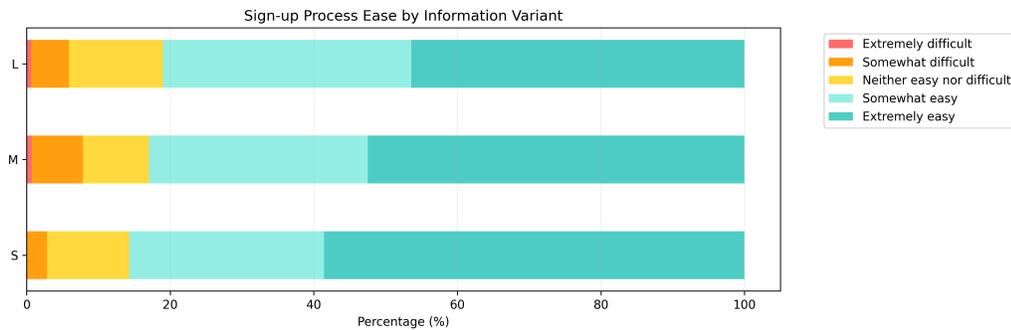


Fig. 11. Q9: Ease of Sign-up Process

When asked about the amount of information provided (Q10), participants reported varied experiences. Most responses indicated that the information was "About what I expected" (S: 43.57%,

M: 36.88%, L: 32.03%) or "Somewhat more than I expected" (S: 32.86%, M: 41.13%, L: 35.29%). A smaller percentage reported that the information was "Much more than expected," with Variant L showing the highest proportion (26.14%) compared to S (14.29%) and M (17.73%). A Chi-square test revealed no significant differences across groups ($\chi^2 = 10.88, p = 0.2086$).

Regarding preferences for settings displayed per page (Q13), the majority preferred "2-3 settings per page" (S: 50.71%, M: 43.26%, L: 45.10%). Participants in Variant M showed a higher preference for "1 setting per page" (38.30%), while participants in Variant L slightly preferred "All settings on a single page" (26.80%). No significant differences were observed ($\chi^2 = 10.29, p = 0.1128$), indicating a general preference for moderate grouping across all variants.

Lastly, participants were asked to evaluate the clarity of the information (Q19). Most found the information "Somewhat clear" (S: 44.29%, M: 45.39%, L: 43.79%). Responses of "Very clear" were more frequent in Variants M (31.91%) and L (29.41%) compared to S (18.57%). A Kruskal-Wallis test revealed statistically significant differences ($H = 8.57, p = 0.0137$), with post hoc analysis showing a significant difference between Variants S and M ($p = 0.0052$). These results suggest that Variant M provided the clearest presentation of information, while Variant S was less effective.

Overall, participants expressed satisfaction with the amount of information (Q20), with most selecting "Somewhat satisfied" (S: 42.14%, M: 43.97%, L: 41.18%) or "Very satisfied" (S: 18.57%, M: 26.24%, L: 23.53%). Dissatisfaction was minimal across all groups, and no significant differences were found ($H = 3.28, p = 0.1937$). These findings indicate that consent flow length does not significantly affect satisfaction levels.

4.1.2 Impact on Comprehension. To assess participants' comprehension and retention of the consent flow content, we asked four questions. For questions requiring single correct answers, participants demonstrated high accuracy. However, for multiple-select questions, participants struggled to identify all correct options, with fewer than 25% of participants selecting all required options correctly. Notably, participants who experienced the long variant (L) showed relatively better overall comprehension of the presented information compared to other variants.

The hardest one being Q15, "What are the features "Now AI" provides? (Select all that apply)". The correct answer is "Smart Assistant,Caption Generation,Image Creation". The vast majority of participants failed to correctly identify all available features. Across different device sizes, the error rates were substantial: 90.71% of users on small variants(S), 87.23% on medium variants (M), and 79.74% on large variants (L) selected incorrect combinations of features. In other words, participants consistently struggled to select the complete and accurate set of features that were actually available across all variants. Variant L had the highest percentage of participants selecting more features. A Chi-square test showed weak evidence of differences across groups ($\chi^2 = 5.19,$

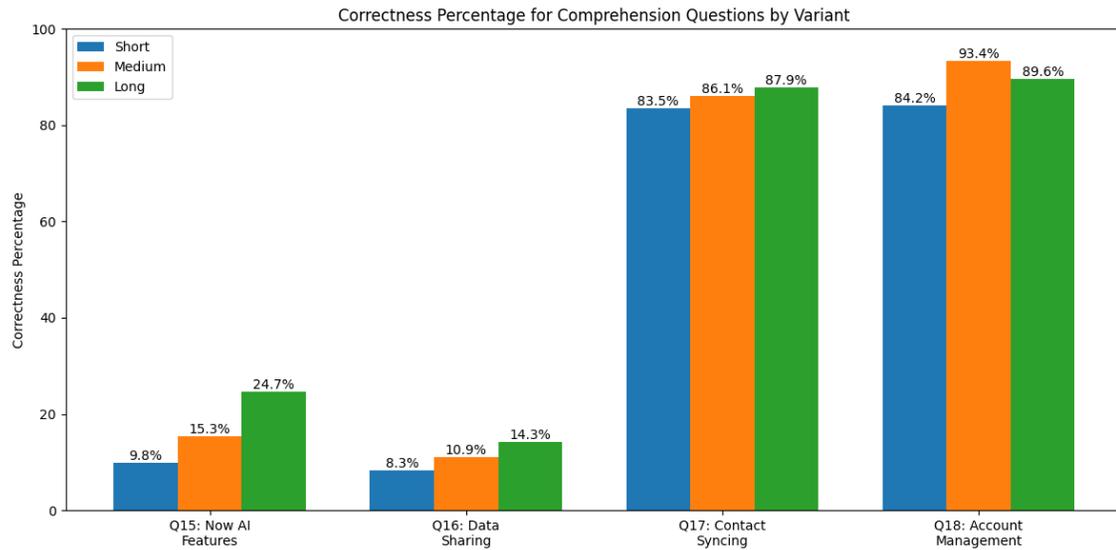


Fig. 12. The correctness of answers to 4 comprehension questions (Q15-Q18) by group. Participants in the long variants (L) displayed a better comprehension of the information they are presented.

$p = 0.0745$). This suggests that participants in variant L may have been slightly more attentive or informed about the features.

Another question: "Which of the following data are being shared if you select the option to share basic account details and BuyNow purchases? (Select all that apply)" (Q16). The correct answer is "Purchase History, Search Preferences". Most of the participants did not select all the correct options, and the majority of each group selected fewer data points (S: 90.71%, M: 90.07%, L: 86.27%). A Chi-square test revealed no significant differences between groups ($\chi^2 = 1.18$, $p = 0.5549$). These findings indicate a similar level of comprehension across all variants.

There are two easier ones, "When can you change the setting related to contact syncing?" (Q17). The correct answer is "Anytime in your account preferences". Most correctly identified that this setting could be changed at any time (S: 84.29%, M: 82.27%, L: 88.89%). No significant differences ($\chi^2 = 1.82$, $p = 0.4024$) were found.

"How can you manage the accounts if you want to unlink them in the future?" (Q18). The correct answer is "Through the Account Management Hub". Most of the participants answered correctly, with variant M having the highest percentage of correct responses (S: 84.29%, M: 94.33%, L: 89.54%). Weak evidence of differences ($\chi^2 = 5.32$, $p = 0.0701$) was found, indicating that variant M may have been slightly more effective in conveying this information.

4.1.3 Impact on Behavior. Behavioral analysis was conducted on survey questions and participant app interaction logs. The results show no differences among the groups.

We first asked participants whether they feel that the information provided on each screen helped them make meaningful choices (Q12). Most the participants felt that the information was sufficient to make meaningful choices, with 69.29% (S), 78.01% (M) and 73.86% (L) agreeing that most or all screens provided enough information. No significant differences between groups ($\chi^2 = 2.28, p = 0.6848$) were found. These results suggest that all variants provided adequate support for decision-making during the sign-up process.

The app interaction logs were analyzed, with 31.43% of the participants in variant S expanded the accordion to show more information about AI features, and 38.30% from M. There are no significant differences ($\chi^2 = 1.17, p = 0.2788$) between them. Furthermore, 50.00% from S and 60.28% from M expanded the accordion to show more BuyNow information. Similarly, no significance ($\chi^2 = 2.60, p = 0.1067$) was found.

4.2 RQ2: What information would users like to have as part of the consent flow, and how does having or not having this information affect their sentiment, comprehension, and behavior?

Our analysis focused on understanding users' information needs and preferences across different variants during the consent flow process. We examined participants' responses regarding information adequacy, desired additional information, and overall satisfaction with the consent options presented.

4.2.1 Information Adequacy and Comprehension. In Q12, participants evaluated whether the screens provided sufficient information for making meaningful choices. The majority across all variants agreed that "Most or all screens provided enough information" (S: 69.3%, M: 78.0%, L: 73.9%). A smaller group indicated that "Some screens provided enough information but others lacked necessary information" (S: 27.1%, M: 19.9%, L: 22.2%), while very few reported that "Most or all screens lacked information" (S: 3.6%, M: 2.1%, L: 3.9%). The variants showed no significant differences ($\chi^2 = 2.28, p = 0.685$).

4.2.2 Areas Requiring Additional Information. : In a multi-select question asking which consent screens would benefit from additional information (Q24), participants consistently identified "Data sharing with 3rd Parties" as the primary concern across all variants (S: 49.3%, M: 53.9%, L: 41.8%), followed by "Data usage for AI Training" (S: 37.9%, M: 39.7%, L: 35.9%). "Cross App Data Sharing" received fewer requests (S: 23.6%, M: 19.1%, L: 19.6%), while a substantial portion indicated no need for additional information (S: 32.1%, M: 30.5%, L: 41.2%). The chi-square test showed no significant differences ($\chi^2 = 6.07, p = 0.415$).

Qualitative insights from follow-up responses revealed specific information needs, including "Third Parties Involved," "Types of Data Being Shared," "Opt-out Policies," and "Third Party Usage Details." Notably, 20% of participants emphasized wanting more information about data usage purposes and requested concrete examples. One participant captured this sentiment: *"Anything and everything. What is being collected, how it is used, if it is sold, how long it is retained for. It is my data, and I deserve to know as much as possible about it."* This quote highlights the depth of information some users expect in consent flows. These themes are visualized in Figure 13.

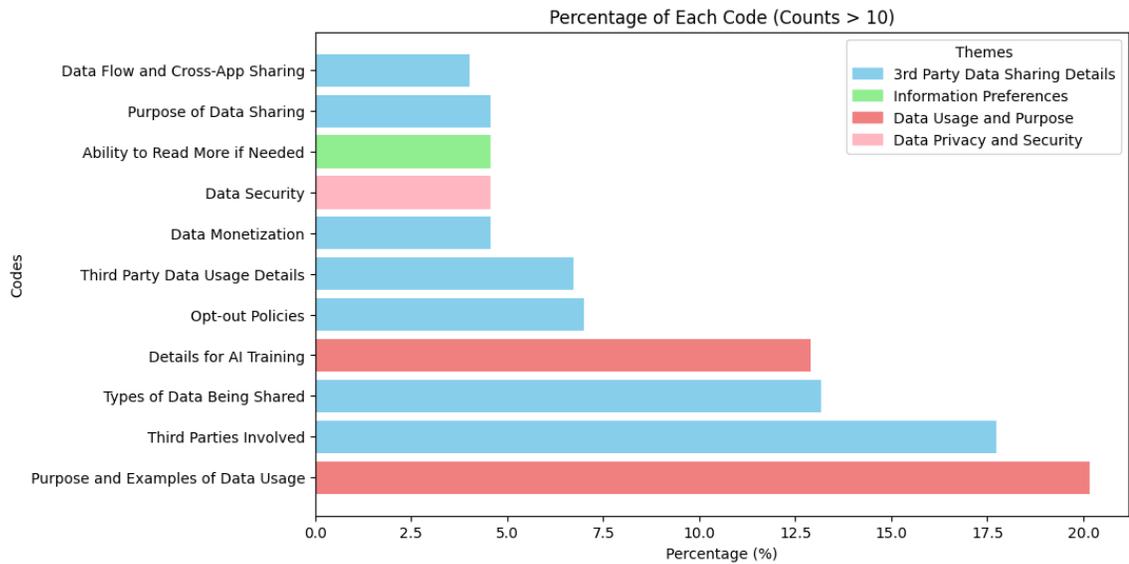


Fig. 13. Participant's coded free response to the question "What kind of information did you expect to see?" on the Data Sharing with 3rd Parties, Cross-App Data Sharing, and Introducing AI screens

4.2.3 User Satisfaction with Options and Information. In Q21, most participants expressed satisfaction with the current options available (S: 55.0%, M: 63.8%, L: 62.1%). Among those who desired more options, preferences varied by variants, with "AI Training" options being more requested on small variants (30.0%) and "Cross App Data Sharing" slightly more desired on large variants (19.6%). "Data sharing with third parties" was consistently desired across all variants (S: 23.6%, M: 22.0%, L: 22.2%). The chi-square test showed no significant differences across variants ($\chi^2 = 3.45$, $p = 0.750$).

Finally, the Q22 responses indicated that the majority found the provided information sufficient (S: 60.0%, M: 68.1%, L: 68.0%), while a smaller group desired more information (S: 34.3%, M: 24.8%, L: 23.5%), and very few preferred less information (S: 5.7%, M: 7.1%, L: 8.5%). The chi-square test did not reveal significant differences between variants ($\chi^2 = 3.71$, $p = 0.447$).

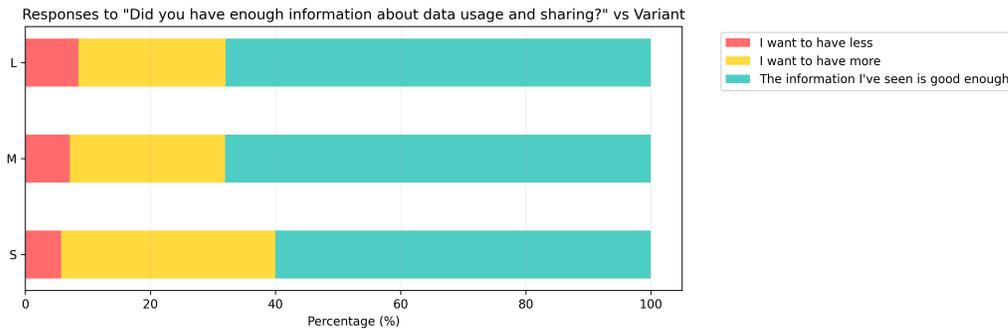


Fig. 14. Q22: Had Enough Information From the Consent Flow

4.3 RQ3: How do different levels of control over data sharing affect users' willingness to engage with AI features in a social media platform?

4.3.1 *Level of Control and Perceived Control.* To assess how users felt about their control over data, we analyzed their responses to two key questions.

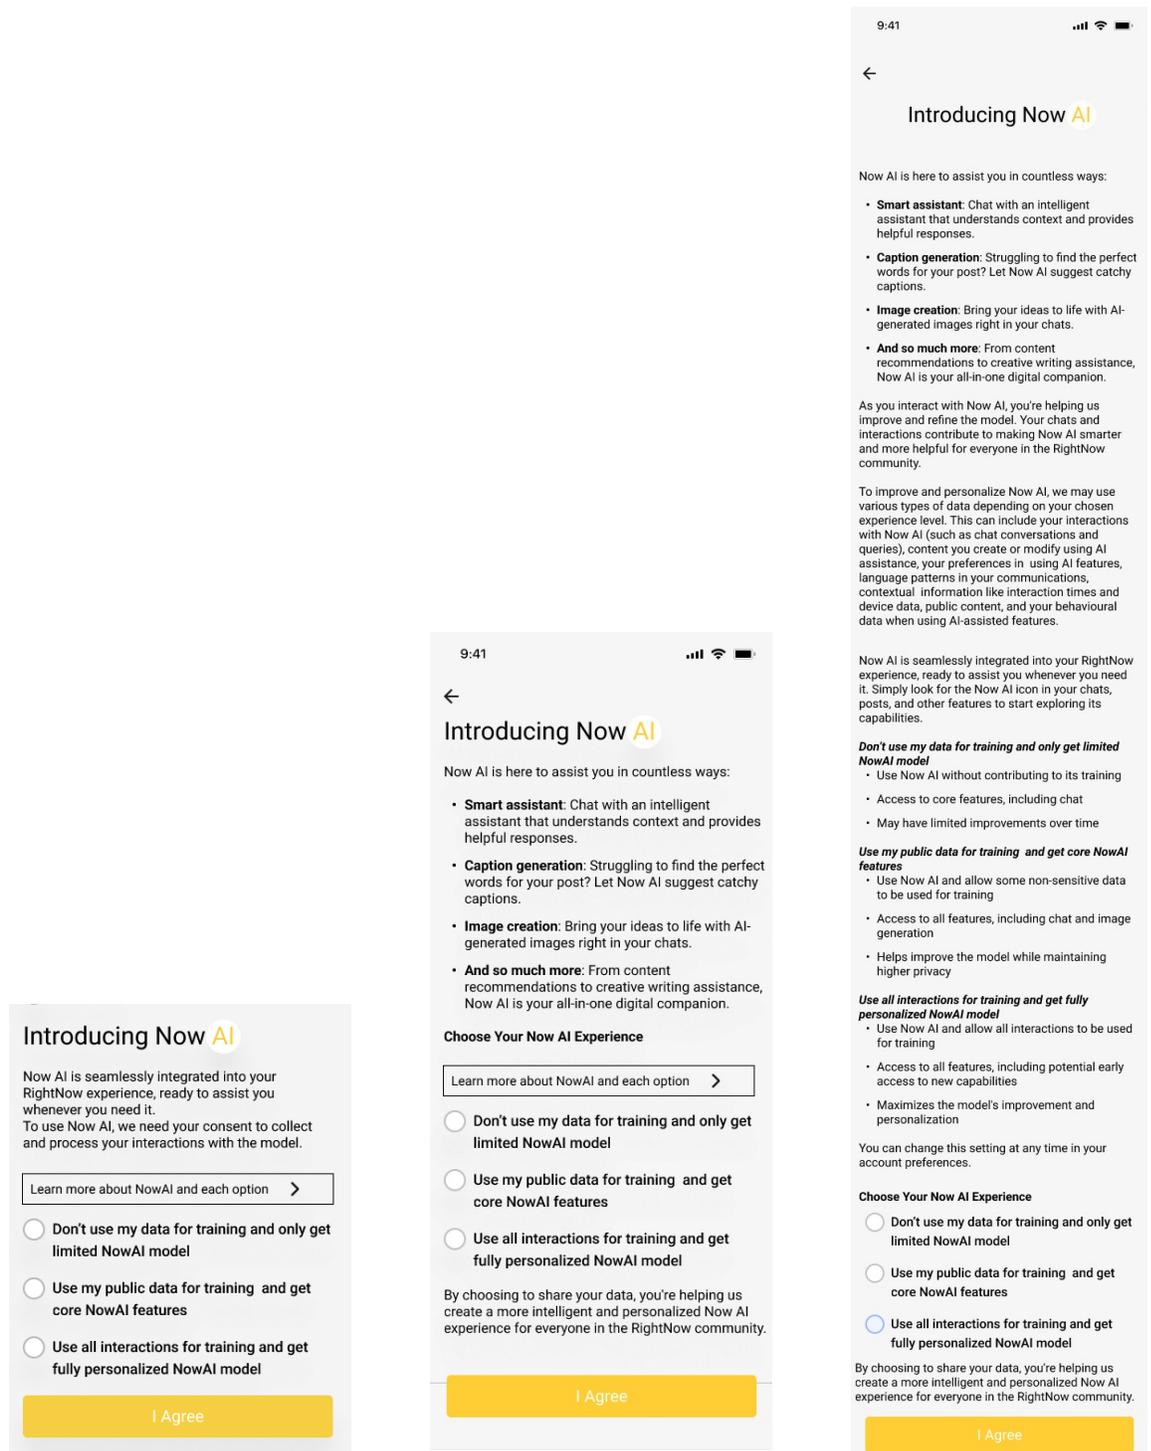
Q11 asked participants about their perceived level of control over their data while using the RightNow app. The results indicated that the majority of participants across all conditions felt they had some level of control over their data (41.4% in variant 1, 52.3% in variant 2, 46.9% in variant 3). However, participants in the granular control variants (S3, M3, L3) reported slightly higher levels of control ($\mu = 4.05$, $\sigma = 0.79$) compared to those in the notice-only (S1, M1, L1) and opt-in/out variants (S2, M2, L2), where responses were more varied. These results suggest that granular control over data sharing leads to a higher sense of empowerment among users.

For Q25, which asked whether the option to partially opt-out of data sharing made participants feel more in control, a significant portion of users in the granular control variants (S3: 57.14%, M3: 58.33%, L3: 57.41%) strongly agreed or agreed that this option enhanced their sense of control. In contrast, the notice-only variants (S1: 5.71%, M1: 1.52%, L1: 7.41%) and the opt-in/out variants (S2: 25%, M2: 28%, L2: 26.54%) had notably lower percentages of participants reporting increased control. These findings suggest that the ability to opt-out, particularly with granular control, plays a significant role in enhancing users' sense of control over their data.

4.3.2 *Clarity of Information on Data Sharing and AI.* The clarity of the information provided regarding data sharing and AI features was assessed through Q19, which asked participants how clear they found the information about data sharing and AI features on the screens. The majority of participants across all conditions rated the information as "somewhat clear" or "very clear." However, participants in the granular control conditions (S3, M3, L3) rated the clarity slightly

higher ($\mu = 3.83$, $\sigma = 1.00$), suggesting that more detailed control options also contributed to clearer communication regarding data usage.

4.3.3 Engagement with AI Features. To assess users' willingness to engage with AI features, we analyzed both self-reported engagement and actual app interaction data. Participants in the granular control variants (S3, M3, L3) reported a higher willingness to engage with AI features, with 60% of participants in the granular control group indicating they were likely to engage, compared to 50% in the opt-in/out group and 38% in the notice-only group. This suggests that the ability to control data sharing more specifically enhances users' willingness to use AI features. App interaction logs showed similar patterns. Participants in the granular control conditions interacted more with AI-related settings, such as expanding the accordion to view more information about AI training. For instance, 38.30% of participants in the M3 condition interacted with AI settings, compared to 31.43% in the S3 condition and 28.57% in the L3 condition. Although these differences were not statistically significant, they suggest a trend toward greater engagement with AI features when users have granular control over their data.



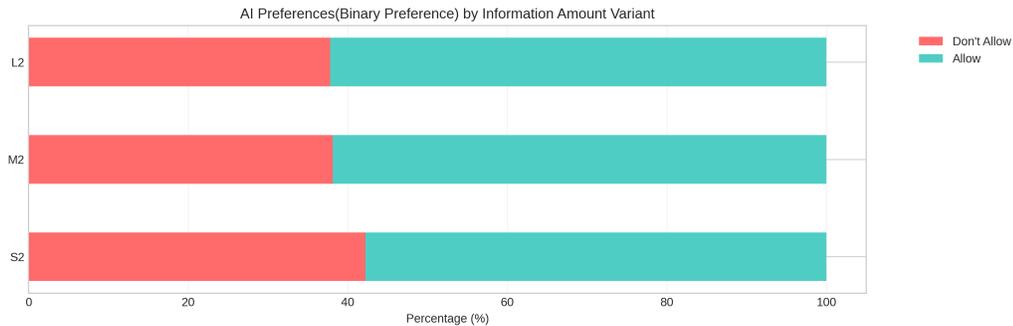
(a) Short-form Granular Controls (S3) - Most information hidden behind accordions

(b) Medium-form Granular Controls (M3) - Balanced visible and hidden information

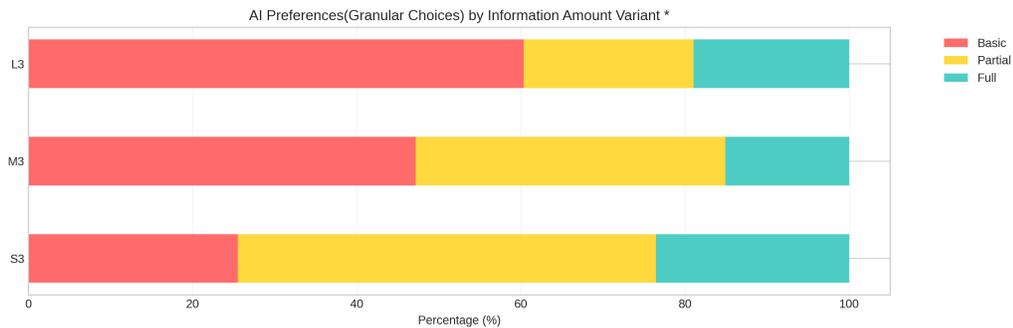
(c) Long-form Granular Controls (L3) - All information immediately visible

Fig. 15. Length variations of the granular control interface (Variant 3) that showed statistically significant differences in user engagement. All variants contain the same comprehensive information but differ in their initial presentation: S3 (left) keeps most information hidden behind expandable accordions, M3 (middle) balances visible and hidden information, and L3 (right) displays all information upfront without requiring user interaction to access details.

The analysis of user preferences across three interface variants (S3, M3, and L3) revealed significant differences in the choices users made regarding AI settings. The variants are shown in Figure 15. The Basic setting was most preferred by users in L3 (60.34%), followed by M3 (47.17%) and S3 (25.49%). On the other hand, S3 users exhibited a strong preference for the Partial setting (50.98%). The statistical analysis, including a chi-square test, showed strong evidence against the null hypothesis ($p = 0.0039$), indicating that the interface design influenced user preferences. Pairwise comparisons further highlighted that L3 users significantly preferred Basic settings over S3 ($p = 0.0013$), while S3 users favored Partial settings ($p = 0.0037$).



(a) Variant 2 (Binary Choice) App Preferences



(b) Variant 3 (Granular Choices) App Preferences

Fig. 16. Percentage of participants in each condition who selected each setting in the RightNow app. The settings with answer choices corresponding to groups of people or types of information appear in the right column. We noted no statistically significant difference in Variant 2 (Binary Choice) preferences, but Variant 3 (Granular Choices) had statistically significant differences.

These findings highlight the influence of interface design on user choices related to AI settings. L3 users tend to prefer more basic settings. The direct and transparent presentation of all the details may make users more inclined to make decisions that minimize data sharing. While S3 users leaned toward greater customization (Partial). The Partial setting allows users to selectively share

additional data, indicating that the compact presentation of choices in this design may encourage users to seek more granular control over what data they share, rather than opting for a more restrictive or more expansive choice. These themes are visualized in Figure 16b.

4.4 Key Findings

Our analysis revealed three key findings addressing our research questions on how consent flow design affects user engagement with privacy interfaces:

- (1) **Content Length and Comprehension.** Users found the medium (M) variant more comprehensible than the short (S) variant, but the long (L) variant was not significantly more comprehensible. Users also prefer the L variant more comprehensible than the others. The two findings suggest that, while detailed information improves clarity, careful consideration must be given to how this information is presented.
- (2) **Information Needs and Transparency.** Our findings revealed that the provided information meets most users' self-reported needs. However, while the information included in our consent flows did not significantly impact user sentiment, comprehension, or behavior, participants consistently expressed a desire for more detailed information about third-party data sharing and AI training practices. This suggests current consent flows may not be addressing users' key privacy concerns.
- (3) **Control Options and AI Engagement.** Among the variants offering granular AI controls (S3, M3, L3), users presented with the longer interface (L3) tended to choose basic privacy settings with minimal data sharing, while those viewing the compact interface (S3) preferred more granular control over their AI data sharing preferences.

5 Discussion

5.1 Implications

Our findings provide practical insights for designing consent flows in modern digital platforms. The results suggest that providing detailed information in consent interfaces slightly improves user comprehension, but the presentation of this information requires careful consideration. While longer formats enhance clarity, designers must focus on effective information organization to prevent cognitive overload. The persistent gap between provided information and user needs highlights the importance of addressing specific user concerns about third-party data sharing and AI training practices in consent flows. Additionally, the influence of interface design on AI privacy choices suggests that designers should consider how different presentation formats affect user decision-making. Longer interfaces appear to encourage more privacy-conscious choices, while compact interfaces promote more granular control. These findings indicate that platform designers

must carefully consider the trade-offs between interface length, information presentation, and the level of control offered to users when developing privacy interfaces.

5.2 Limitations

Our study, while offering valuable insights, presents several notable limitations that warrant acknowledgment. The simulated app environment utilized in our research may not fully encapsulate the complexities of real-world user behavior, necessitating further validation in live systems to effectively generalize our findings. Additionally, our participant recruitment through Prolific potentially resulted in a non-representative sample, suggesting that future research would benefit from expanding recruitment to more diverse demographics. The 12-minute survey format imposed time constraints that may have influenced user decisions, as participants might approach such tasks differently in real-world contexts without such temporal limitations. Furthermore, our findings are specifically contextualized within the tested social media and marketplace app environment, indicating that extending the study to other types of platforms could enhance the broader applicability of our results. Finally, our findings may be specific to the particular interface designs and terminology used in our study; future research should explore alternative design variations and wording choices to establish the broader generalizability of these results.

5.3 Future Research Directions

To build upon our findings, we recommend several key areas for future research. Longitudinal studies would be valuable to examine how consent choices influence long-term user behavior and platform engagement. Additionally, research should investigate the impact of different information presentation methods, particularly focusing on how varying wording affects user trust and decision-making processes. Future studies should also explore the role of detailed explanations of AI features in fostering user adoption and trust in AI-enabled platforms. Furthermore, examining the impact of content length and control mechanisms across different cultural and demographic groups would provide valuable insights into the universality and effectiveness of these approaches.

6 Ethics

We obtained approval from Carnegie Mellon University’s Institutional Review Board (IRB) for our online survey. All participants were presented with an IRB-approved consent form and were informed of their right to withdraw from the study at any time without penalty. Participants were compensated at a rate of \$15/hour, receiving \$3 for completing the 12-minute survey. Given that the RightNow app was purely a simulated app experience, we did not collect any personal information from participants except Prolific IDs, which were used solely for compensation purposes. The study presented minimal risk to participants as all scenarios were simulated.

7 Conclusion

Our study demonstrates that consent flow design significantly influences how users understand and interact with privacy choices, particularly in contexts involving AI features and data sharing controls. The findings reveal that detailed information in consent interfaces enhances user comprehension, challenging the common assumption that simpler interfaces are always better. This presents a key challenge for designers: balancing comprehensive information presentation with effective interface design.

These insights offer specific implications for privacy interface design within our tested contexts. Our results indicate that longer formats enhanced clarity and promoted privacy-conscious decisions among our participants, while compact interfaces appeared to facilitate more granular control over data sharing in our tested scenarios. Additionally, the study revealed a notable gap between provided information and participant needs, specifically regarding third-party data sharing and AI training practices in the tested interfaces. These findings suggest that the studied consent flows may benefit from restructuring to better address these identified privacy concerns while maintaining effective information presentation.

Future work should explore how to effectively incorporate detailed information about data practices while optimizing interface design for different user preferences. Additionally, research could investigate how varying presentation formats influence user decision-making across different demographic groups and cultural contexts, particularly in relation to AI privacy controls and data sharing preferences.

Acknowledgments

We would like to thank the Carnegie Mellon Privacy Engineering Master's program for providing the education and tools necessary to conduct this study. We are grateful to Professor Lorrie Cranor for her continuous support and feedback. We thank Meta for sponsoring this project, and especially Eric Foster Moore, whose time and insights into the social media industry were invaluable to this work. We thank Professor Sarah Scheffler for her thoughtful feedback and the CMU staff for their assistance with funding arrangements and server setup. Finally, we extend our gratitude to all the participants who volunteered their time for our study.

References

- [1] Eman Alashwali, Andrea Miller, Alexandra Nisenoff, Kyerra Norton, Ellie Young, Prerna Bothra, Mandy Lanyon, and Lorrie Faith Cranor. 2024. Interface Design to Support Informed Choices When Users Face Numerous Privacy Decisions. (2024). Unpublished manuscript.
- [2] Travis Breaux. 2024. An Introduction to Privacy for Technology Professionals. <https://iapp.org/resources/article/an-introduction-to-privacy-for-technology-professionals> [Online; accessed 11. Dec. 2024].
- [3] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. 2018. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior* 81 (2018), 42–51.
- [4] Verena Distler, Matthias Fassl, Hana Habib, Katharina Krombholz, Gabriele Lenzini, Carine Lallemand, Vincent Koenig, and Lorrie Faith Cranor. 2023. Empirical Research Methods in Usable Privacy and Security. In *Human Factors in Privacy Research*. Springer, Cham, Switzerland, 29–53. https://doi.org/10.1007/978-3-031-28643-8_3
- [5] Hana Habib and Lorrie Faith Cranor. 2022. Evaluating the Usability of Privacy Choice Mechanisms. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. USENIX Association, Boston, MA, 273–289. <https://www.usenix.org/conference/soups2022/presentation/habib>
- [6] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376511>
- [7] Hana Habib, Sarah Pearman, Ellie Young, Ishika Saxena, Robert Zhang, and Lorrie Faith Cranor. 2022. Identifying User Needs for Advertising Controls on Facebook. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1, Article 59 (April 2022), 42 pages. <https://doi.org/10.1145/3512906>
- [8] Silas Hsu, Kristen Vaccaro, Yin Yue, Aimee Rickman, and Karrie Karahalios. 2020. Awareness, Navigation, and Use of Feed Control Settings Online. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376583>
- [9] Mark J Keith, Courtenay Maynes, Paul Benjamin Lowry, and Jeffrey Babb. 2014. Privacy fatigue: The effect of privacy control complexity on consumer electronic information disclosure. In *International Conference on Information Systems (ICIS 2014)*, Auckland, New Zealand, December. 14–17.
- [10] Yabing Liu, Krishna P Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing facebook privacy settings: user expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. 61–70.
- [11] Michelle Madejski, Maritza Johnson, and Steven M Bellovin. 2012. A study of privacy settings errors in an online social network. In *2012 IEEE international conference on pervasive computing and communications workshops*. IEEE, 340–345.
- [12] Stefania Passera. 2015. Beyond the Wall of Text: How Information Design Can Make Contracts User-Friendly. In *Design, User Experience, and Usability: Users and Interactions*. Springer, Cham, Switzerland, 341–352. https://doi.org/10.1007/978-3-319-20898-5_33
- [13] Katherine Strater and Heather Richter Lipford. 2008. Strategies and struggles with privacy in an online social networking community. In *People and Computers XXII Culture, Creativity, Interaction*. BCS Learning & Development.
- [14] Karl Van Der Schyff, Greg Foster, Karen Renaud, and Stephen Flowerday. 2023. Online privacy fatigue: a scoping review and research agenda. *Future Internet* 15, 5 (2023), 164.

- [15] Nicholas Vollmer. 2023. Recital 32 EU General Data Protection Regulation (EU-GDPR). *SecureDataService* (April 2023). <https://www.privacy-regulation.eu/en/recital-32-GDPR.htm>

A Experimental Design

Table 3. RQ1: How do the number of pages and explanatory information in the consent flow impact user sentiment, comprehension, and behavior?

Sub RQ	Variants (Independent Variables)	Data Sources	Actual Survey Questions (Qx from Qualtrics)
Sentiment	Short vs Medium vs Long	Survey results, Click logs	<p>Q6: How fatigued did you feel during the sign-up process for the RightNow app?</p> <p>Q7: How likely would you have been to complete the process if you were signing up for the RightNow app in real life?</p> <p>Q9: How easy or difficult was the sign-up process for you?</p> <p>Q10: How did you feel about the amount of information provided during the signup process?</p> <p>Q15: During the signup process we asked you to make choices for several settings. How would you prefer to have these choices presented?</p> <p>Q19: How clear was the information provided on the screens regarding data sharing and AI features?</p> <p>Q20: How satisfied were you with the amount of information provided on the screens?</p>

Sub RQ	Variants (Independent Variables)	Data Sources	Actual Survey Questions (Qx from Qualtrics)
Comprehension	Short vs Medium vs Long	Survey results, App click logs	Q15: During the signup process we asked you to make choices for setup settings. How would you prefer to have these choices presented? Q15.5: What are the features “Now AI” provides? (Select all that apply) Q16: Which of the following data are being shared if you select the option to share basic account details and BuyNow purchases? (Select all that apply) Q17: When can you change the setting related to contact syncing? Q18: How can you manage the accounts if you want to unlink them in the future? Q22: Did you have enough information about data usage and sharing?
Behavior	Short vs Medium vs Long	App click logs, Time spent	Q12: Did you feel the information provided on each screen assisted you in making meaningful choices?

Table 4. RQ2: What information would users like to have as part of the consent flow, and how would having or not having this information affect their sentiment, comprehension, and behavior?

Sub RQ	Variants (In-dependent Variables)	Data Sources	Feedback
Sentiment	Information level (basic, moderate, detailed)	Survey results	Q12: Did you feel the information provided on each screen assisted you in making meaningful choices? Q21: Would you like more options on any screen? Q22: Did you have enough information about data usage and sharing?
Comprehension	Information level (basic, moderate, detailed)	Survey results	Q24: Which of these settings, if any, did you want more information about before making a selection? Q24.1: For the question 'Which of these settings, if any, did you want more information about before making a selection?', you selected: in Q24 What kind of information did you expect to see?
Behavior	Information level (basic, moderate, detailed)	App click logs, Survey results	Q12: Did you feel the information provided on each screen assisted you in making meaningful choices?

Table 5. RQ3: How do different levels of control over data sharing affect users’ willingness to engage with AI features in a social media platform?

Sub RQ	Variants (Independent Variables)	Data Sources	Feedback
-	Short 1 vs Short 2 vs Short 3, Medium 1 vs Medium 2 vs Medium 3, Long 1 vs Long 2 vs Long 3	Survey results, App click logs	Q13: How do you feel about the level of control you have over your data when using the RightNow app? Q25: Did having an option to partially opt-out of data sharing (e.g., AI Training, third-party advertising) make you feel more in control? Q19: How clear was the information provided on the screens regarding data sharing and AI features?

B Additional Tables

Table 6. Participant demographic information for age, gender, and race

Age	Years
Minimum Age Range	18-24
Maximum Age Range	65 or older
Average Age Range	25-34
Gender	% Participants
Male	51.84%
Female	46.77%
Non-binary	1.15%
Prefer not to answer	<1%
Prefer to self-describe	<1%
Race/Ethnic Identity (non-exclusive)	Participant Count
White	253
Black or African American	80
Asian	36
Hispanic and/or Latino/Latina/Latinx	20
American Indian or Alaskan Native	1
Native Hawaiian or Pacific Islander	1
Other	5

Table 7. Participant demographic information for highest level of education, household income, and computer-field experience.

Highest education level	% Participants
Less than high school diploma	2.3%
High school graduate (high school diploma or equivalent including GED)	10.6%
Some college but no degree	27.19%
Associate degree in college	11.52%
Bachelor's degree	34.56%
1 Master's degree	11.06%
Doctoral degree	0.92%
Professional degree (JD, MD)	1.84%
Other	0%
Have a computer-related field of education or employment	% Participants
Yes	35.25%
No	63.13%
Household income	% Participants
Less than \$40,000	26.96%
\$40,000 to \$59,999	18.89%
\$60,000 to \$79,999	17.05%
\$80,000 to \$99,999	11.75%
\$100,000 to \$149,999	13.36%
\$150,000 or more	6.91%
Prefer not to answer	3.46%

Table 8. Participant social media and online marketplace app use

Social media use daily frequency	% Participants
Multiple times per day	46.31%
Multiple times per hour	19.59%
At least once an hour	19.35%
At least once a day	11.98%
Less than once a day	2.53%
Social media apps used during the past week	% Participants
Discord	36.41%
BeReal	1.15%
Facebook	70.28%
Instagram	73.96%
LinkedIn	29.03%
SnapChat	32.72%
Threads	8.29%
TikTok	48.39%
Twitter/X	49.77%
WeChat	0.92%
Other	19.59%
None	0%
Marketplace app use daily frequency	% Participants
Multiple times per day	8.53%
Multiple times per hour	1.84%
At least once an hour	1.38%
At least once a day	25.81%
Less than once a day	60.83%
Marketplace apps used during the past week	% Participants
Amazon	87.79%
eBay	32.95%
Facebook Marketplace	30.18%
Etsy	21.66%
Craigslist	10.14%
OfferUp	3.92%
I haven't used any online marketplace apps this week	6.91%
Other (Please describe)	4.38%

Table 9. Participant distribution by variant

Condition	# Participants
Short 1 AI Training Notice	44
Short 2 AI Opt In/Opt Out	45
Short 3 AI Granlar Opt Out	51
Medium 1 AI Training Notice	46
Medium 2 AI Opt In/Opt Out	42
Medium 3 AI Granlar Opt Out	53
Long 1 AI Training Notice	50
Long 2 AI Opt In/Opt Out	45
Long 3 AI Granlar Opt Out	58
Total surveys analyzed	434

Table 10. Codebook for free response to the question "What kind of information did you expect to see?" on the Data Sharing with 3rd Parties, Cross-App Data Sharing, and Introducing AI screens

Code	Description
Types of Data Being Shared	Users wanted to know what specific data points were being shared.
Third Parties Involved	Users wanted to know who the third parties receiving the data were, including examples or names of companies.
Purpose of Data Sharing	Users wanted to understand why the data was being shared and for what specific purposes.
Data Flow and Cross-App Sharing	Users were concerned about how data was shared across different apps, including cross-platform sharing.
Third Party Data Usage Details	Users wanted specific information on how third parties would use their data, including if it would be sold or shared further.
Data Monetization	Users were curious if their data would be sold to third parties and how it might be monetized.
Opt-Out Policies	Desire for an option to completely opt out of data sharing, including questions about "Do Not Sell" policies.
Purpose and Examples of Data Usage	Requests for an explanation of why specific data is collected, including examples of how data is used.
Details for AI Training	Desire for more specific information on what data is used for AI training.
Data Security	Questions related to the security measures taken to protect data.
Ability to Read More if Needed	Preference for the option to access more detailed information through a link.

Table 11. Codebook for free response to the question "What factors would have influenced your decision to continue or quit in real life?"

Code	Description
Ease of Use	Accessibility and efficiency of the application's onboarding process and registration flow
User Interface	Design elements, layout, and overall visual organization of the application's interface
Readability	Clarity and comprehensibility of textual content and information presentation
Screen Length	Total number of interface screens in the onboarding sequence
Text Density	Distribution and volume of textual content per screen
Decision Overload	Cognitive burden from excessive sequential decision points
Privacy Concerns	User apprehensions regarding data protection and information security
User Control	Availability and accessibility of opt-out mechanisms and user preferences
Data Collection	Scope and nature of personal information required during user registration
Data Sharing	Protocols and policies for third-party data distribution
User Benefits	Perceived advantages, incentives, and value propositions offered to users
Utility	Practical application and functional relevance to user needs
Brand Trust	User confidence and reliability perception in the service provider
AI Skepticism	User reservations and concerns regarding artificial intelligence implementation
Purchase Visibility	Social media integration of transaction-related information
Social Influence	Impact of peer and family recommendations on adoption
Time Investment	Duration required for complete registration process
Technical Performance	Application's operational efficiency and functionality
System Permissions	Required access authorizations and system-level privileges

Table 12. Codebook for free response to the question "Please tell us what made you feel that the app was encouraging you to give up your privacy."

Code	Description
Data Usage and Users	User concerns regarding data utilization patterns and the entities accessing their personal information
Data Sharing Across Apps and Third Parties	Apprehensions about cross-platform data distribution and third-party information access protocols
Lack of Choice in Data Sharing	Perceived constraints in data sharing options and concerns about predetermined sharing configurations
Interface Design Encouraging Data Sharing	User interface elements and patterns that potentially influence or direct users toward expanded data sharing
Lack of Transparency	Insufficient clarity and communication regarding intended data usage and processing methods
AI-Related Data Concerns	Specific concerns regarding data utilization in artificial intelligence systems and machine learning training processes
App Utility and Privacy Trade-off	Perceived compromise between application functionality and required levels of data disclosure

Table 13. Statistical Analysis of User Responses Across All Questions (Part 1)

Question	Options	Variants (%)			Test	Statistic	p-value
		S	M	L			
Q6. How fatigued did you feel during the sign-up process for the RightNow app?	Moderately fatigued Not fatigued at all Quite fatigued Slightly fatigued Very fatigued	11.43 58.57 3.57 26.43 0.00	10.64 58.16 2.84 27.66 0.71	12.42 51.63 4.58 28.76 2.61	Kruskal-Wallis	2.520	0.284
Q7. How likely would you have been to complete the process if you were signing up for the RightNow app in real life?	Definitely would complete Definitely would quit Likely would complete Likely would quit Neutral	32.86 4.29 45.71 5.71 11.43	27.66 4.96 40.43 12.06 14.89	24.18 4.58 41.83 11.11 18.30	Kruskal-Wallis	5.820	0.054
Q9. How easy or difficult was the sign-up process for you?	Extremely difficult Extremely easy Neither easy nor difficult Somewhat difficult Somewhat easy	0.00 58.57 11.43 2.86 27.14	0.71 52.48 9.22 7.09 30.50	0.65 46.41 13.07 5.23 34.64	Kruskal-Wallis	4.227	0.121
Q10. How did you feel about the amount of information provided during the signup process?	About what expected Much less than expected Much more than expected Somewhat less than expected Somewhat more than expected	43.57 2.14 14.29 7.14 32.86	36.88 2.84 17.73 1.42 41.13	32.03 1.31 26.14 5.23 35.29	Chi-square	10.880	0.209
Q11. How do you feel about the level of control you have over your data when using the RightNow app?	A lot of control Little control Neutral Some control Very little control	18.57 13.57 14.29 49.29 4.29	26.24 10.64 9.93 48.23 4.96	24.18 8.50 18.30 43.14 5.88	Kruskal-Wallis	2.333	0.311
Q12. Did you feel the information provided on each screen assisted you in making meaningful choices?	Most/all lacked info Most/all had enough info Some had enough info	3.57 69.29 27.14	2.13 78.01 19.86	3.92 73.86 22.22	Chi-square	2.278	0.685
Q13. During the signup process we asked you to make choices for several settings. How would you prefer these presented?	1 setting per page (6 pages) 2-3 settings per page All 6 settings on single page Other	20.71 50.71 26.43 2.14	38.30 43.26 18.44 0.00	26.14 45.10 26.80 1.96	Chi-square	10.293	0.113
Q15. What are the features "Now AI" provides?	Correct Incorrect	9.29 90.71	12.77 87.23	20.26 79.74	Chi-square	5.194	0.075
Q16. Which data is shared in basic account details?	Correct Incorrect	9.29 90.71	9.93 90.07	13.73 86.27	Chi-square	1.178	0.555
Q17. When can you change contact syncing settings?	Correct Incorrect	84.29 15.71	82.27 17.73	88.89 11.11	Chi-square	1.821	0.402
Q18. How can you manage account unlinking?	Correct Incorrect	84.29 15.71	94.33 5.67	89.54 10.46	Chi-square	5.316	0.070
Q19. How clear was the information provided on screens regarding data sharing and AI features?*	Neutral Somewhat clear Somewhat unclear Very clear Very unclear	25.00 44.29 10.00 18.57 2.14	9.93 45.39 10.64 31.91 2.13	14.38 43.79 10.46 29.41 1.96	Kruskal-Wallis	8.574	0.014
Q20. How satisfied were you with the amount of information provided on the screens?	Neither satisfied nor dissatisfied Somewhat dissatisfied Somewhat satisfied Very dissatisfied Very satisfied	23.57 12.86 42.14 2.86 18.57	16.31 11.35 43.97 2.13 26.24	22.88 8.50 41.18 3.92 23.53	Kruskal-Wallis	3.283	0.194

Table 14. Statistical Analysis of User Responses Across All Questions (Part 2)

Question	Options	Variants (%)			Test	Statistic	p-value
		S	M	L			
Q21. Would you like more options on any screen? (select all that apply)	Cross App Data Sharing	15.71	17.73	19.61	Chi-square	3.452	0.750
	Data sharing with third parties	23.57	21.99	22.22			
	Data usage for AI Training	30.00	24.82	22.22			
	No, I was satisfied	55.00	63.83	62.09			
Q22. Did you have enough information about data usage and sharing?	I want to have less	5.71	7.09	8.50	Chi-square	3.707	0.447
	I want to have more	34.29	24.82	23.53			
	Information is good enough	60.00	68.09	67.97			
Q23. Which settings made you concerned about privacy? (select all that apply)	Cross App Data Sharing	30.71	26.95	33.99	Chi-square	1.979	0.922
	Data sharing with 3rd Parties	77.86	75.18	74.51			
	Data usage for AI Training	41.43	47.52	47.06			
	None	13.57	13.48	15.69			
Q24. Which settings did you want more information about? (select all that apply)	Cross App Data Sharing	23.57	19.15	19.61	Chi-square	6.073	0.415
	Data sharing with 3rd Parties	49.29	53.90	41.83			
	Data usage for AI Training	37.86	39.72	35.95			
	None	32.14	30.50	41.18			
Q25. Did partial opt-out options make you feel more in control?	Agree	55.00	56.74	60.78	Kruskal-Wallis	1.146	0.564
	Disagree	3.57	3.55	3.92			
	Neutral	12.14	10.64	10.46			
	Strongly agree	28.57	28.37	22.88			
	Strongly disagree	0.71	0.71	1.96			
Q26. Did RightNow app settings encourage you to give up privacy?	Most/all encouraged privacy loss	13.57	12.06	15.03	Chi-square	1.485	0.829
	None encouraged privacy loss	48.57	48.23	52.94			
	Some encouraged privacy loss	37.86	39.72	32.03			
Q27. How similar was your approach compared to real life installation?	Less time in real life	9.29	4.96	10.46	Chi-square	3.356	0.500
	More time in real life	19.29	16.31	20.92			
	Similar approach	71.43	78.72	68.63			
Q28. How do you typically select privacy settings for social media apps?	Carefully consider all choices	38.57	47.52	44.44	Chi-square	6.854	0.552
	Choose best as quickly as possible	50.00	46.10	46.41			
	Make random choices	0.71	0.00	1.31			
	Take default settings	7.86	2.84	7.19			
	Other	2.86	3.55	0.65			

Table 15. Analysis of User Responses of Survey Questions for Level of Control Variants for AI

Question	Options	Variants (%)			Test	Statistic	p-value
		1	2	3			
Q11. How do you feel about the level of control you have over your data when using the RightNow app?	A lot of control over my data	17.86	21.97	28.40	Kruskal-Wallis	9.214	0.010*
	Little control over my data	15.00	8.33	9.26			
	Neutral	20.00	15.91	8.02			
	Some control over my data	41.43	52.27	46.91			
	Very little control over my data	5.71	1.52	7.41			
Q19. How clear was the information provided on screens regarding data sharing and AI features?	Neutral	20.71	15.15	13.58	Kruskal-Wallis	2.241	0.326
	Somewhat clear	40.71	46.21	46.30			
	Somewhat unclear	10.71	7.58	12.35			
	Very clear	27.14	29.55	24.07			
	Very unclear	0.71	1.52	3.70			
Q25. Did having an option to partially opt-out of data sharing make you feel more in control?	Agree	57.14	58.33	57.41	Kruskal-Wallis	0.991	0.609
	Disagree	5.71	1.52	3.70			
	Neutral	10.71	11.36	11.11			
	Strongly agree	25.00	28.03	26.54			
	Strongly disagree	1.43	0.76	1.23			

Significant at $p < 0.05$. For Q11, pairwise comparisons showed significant differences between variants 1-2 ($p = 0.013$) and 1-3 ($p = 0.007$).

C App Interaction Visualizations

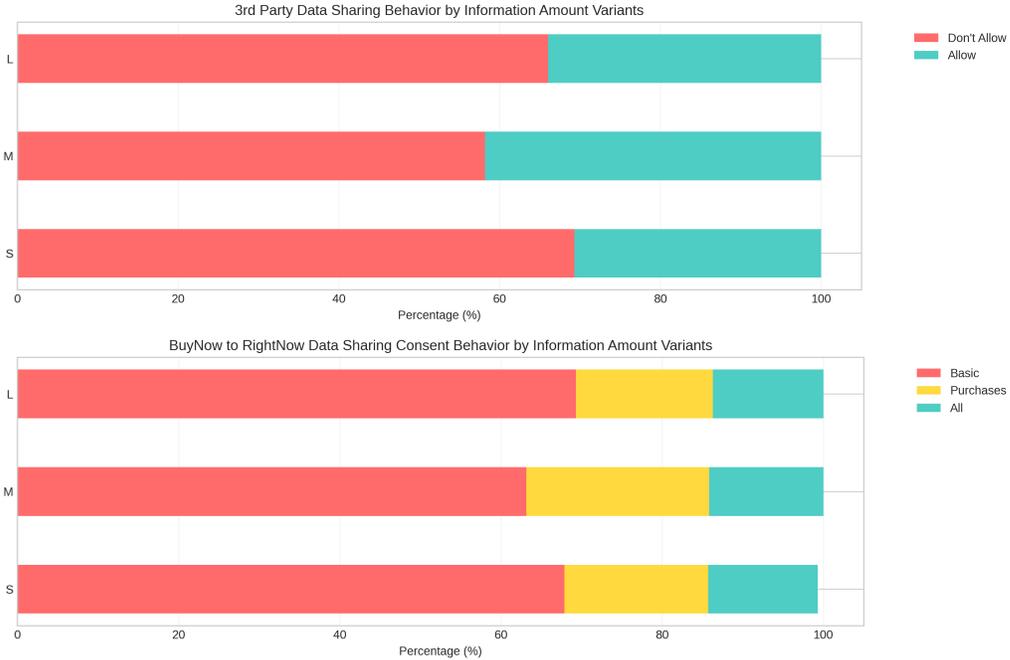


Fig. 17. Percentage of participants in each condition who selected each setting in the RightNow app. The settings with answer choices corresponding to groups of people or types of information appear in the right column. No statistically significant differences we noted between conditions.

D Survey Questions

The full survey questionnaire used in this study is provided below.

Survey Design

In what follows we provide the survey text. Section headings and text between [square brackets] were not visible to participants. We included them here for clarity.

1) Introduction

This survey will take about 12 minutes, including a simulated app onboarding (2-4 minutes), and a main survey (4-6 minutes). You will answer questions about your social media and marketplace app use, and you will view designs for an app and be asked questions about those. For the best experience, please take this survey on your computer screen.

2) Prolific ID & Social Media Use

Q1: Please enter your Prolific ID:

Answer: [free response field]

Please answer these first few questions about your social media use.

Q2: Which of the following social media apps have you used this week? [participants could select multiple options]

Answers: BeReal / Discord / Facebook / Instagram / LinkedIn / Snapchat / Threads / TikTok / X/Twitter / WeChat / I haven't used any social media apps this week [*exclusive] / Other (Please describe): [free response field]

Q3: Approximately, how often do you go on social media throughout the day?

Answers: Multiple times an hour / At least once an hour / Multiple times a day / At least once a day / Less than once a day / Other (Please describe): [free response field]

Q4: Which of the following online marketplace apps have you used this week? (Select all that apply)

Answers: Amazon / eBay / Facebook Marketplace / Etsy / Craigslist / OfferUp / Other (Please describe): [Free response field]

Q5: Approximately how often do you use online marketplace apps?

Answers: a few times per week / a few times per month / a few times per year / once a year or less / never / Other (Please describe): [free response field]

3) Cross-App Onboarding Task Instructions

Imagine you have been using BuyNow, an online marketplace app, for the past few months. The company behind BuyNow is now launching RightNow, a new social media app where you can

connect with friends and family by sharing posts, pictures, videos, and even your purchases from BuyNow!

You will now be looking at a prototype of a new social media app- RightNow. We are especially interested in the sign-up process for a new user of the app.

You will go through the steps as if you had just installed the app. This app uses your BuyNow account username and password to create this new RightNow account. On each screen, please select the settings you think you would choose **if you were signing up for this app in real life**. When you get to the end of the sign-up process you will need to click on "Return to survey".

Note: This simulation is for research purposes only. None of your personal information will be collected during this onboarding task.

Click on the link below to begin.

[Participants were randomly assigned to a condition and completed the signup process in an interactive interface before progressing further in the survey.]

[This is the link to the interface design. Participants would be randomly assigned to one of nine variants below:

<https://www.figma.com/proto/JbnvplLJcOtlrtVQlvS1bu/BuyNow%2FRightNow?page-id=0%3A1&node-id=319-844&node-type=canvas&viewport=-171%2C-850%2C0.19&t=omn0YoJPWxwU0PL7-1&scaling=min-zoom&content-scaling=fixed&starting-point-node-id=319%3A844&show-prot-o-sidebar=1>]

Q: Please confirm that you have clicked on the link and read the webpage.

Answer: I have clicked and read

4) General Sentiment

Please answer the following questions about your overall impression of the RightNow App sign-up process, considering your existing use of BuyNow.

Answer these questions based on your memory.

Note: Please refrain from going back to the app

Q6: How fatigued did you feel during the sign-up process for the RightNow app?

Answers: Not fatigued at all / Slightly fatigued / Moderately fatigued / Quite fatigued / Very fatigued

[shown if participant answered other than Not fatigued in Q8]

Q6.1: Which of the following contributed to your fatigue during the sign-up process (check all that apply)?

Answers: Process had many steps / a lot of information was presented / process was difficult to understand / doesn't care about the agreements and terms / I was tired beforehand

Q7: How likely would you have been to complete the process if you were signing up for the RightNow app in real life?

Answers: Definitely would have completed the process / Likely would have completed the process / Neutral (neither likely to complete nor quit) / Likely would have quit the process / Definitely would have quit the process

Q8: What factors would have influenced your decision to continue or quit in real life?

Answers: [free response field]

Q9: How easy or difficult was the sign-up process for you?

Answers: Extremely easy / Somewhat easy / Neither easy nor difficult / Somewhat difficult / Extremely difficult

Q10: How did you feel about the amount of information provided during the signup process?

Answers: Much shorter than I expected / Somewhat shorter than I expected / About what I expected / Somewhat longer than I expected / Much longer than expected

Q11: How do you feel about the level of control you have over your data when using the RightNow app?

Answers: A lot of control over my data / Some control over my data / Neutral (neither a lot nor little control) / Little control over my data / Very little control over my data

Q12: Did you feel the information provided on each screen assisted you in making meaningful choices?

Answers: Most or all screens provided enough information to make meaningful choices / Some screens provided enough information to make meaningful choices but others lacked necessary information / Most or all screens lacked information needed to make meaningful choices

Q13: During the signup process we asked you to make choices for setup settings. How would you prefer to have these choices presented?

Answers: 1 setting per page / 2-3 settings per page / All settings on a single page (1 page) / Other (Please describe): [free response field]

Q14: For the question '*During the signup process we asked you to make choices for several settings. How would you prefer to have these choices presented?*' you selected:

[participant's response]

Please explain why you chose that answer.

Answer: [free response field]

5) General Comprehension

Q15: What are the features "Now AI" provides? (Choose all that apply)

Answers: Smart Assistant / Video Editing / Caption Generation / Posts Tagging / Image Creation / Posts Music Generation

Q16: Which of the following data are being shared if you select the option to share basic account details and BuyNow purchases? (Select all that apply)

Answers: Purchase History / Search Preferences / Post Interactions / Marketplace Recommendations / Location Data

Q17: When can you change the setting related to contact syncing?

Answers: After 30 days / Only after you complete account setup / Anytime in your account preferences / You cannot change this setting once activated

Q18: How can you manage the accounts if you want to unlink them in the future?

Answers: You cannot unlink them once linked / Through the Account Management Hub / By reinstalling the app / None of the above

6) Data Sharing Related

[For each of the main areas: Cross App Data Sharing, Third Party Advertising, and AI Training]

Q19: How clear was the information provided on the screens regarding data sharing and AI features?

Answers: Very clear / Somewhat clear / Neutral (neither clear nor unclear) / Somewhat unclear / Very unclear

Q20: How satisfied were you with the amount of information provided on the screens?

Answers: Very satisfied / Somewhat satisfied / Neither satisfied nor dissatisfied / Somewhat dissatisfied / Very unsatisfied

Q21: Would you like more options on any screen? (select all that apply)

Answers: Cross App Data Sharing (BuyNow to RightNow) / Data sharing with third parties / Data usage for AI Training / No, I was satisfied

Q22: Did you have enough information about data usage and sharing?

Answers: I want to have more / I want to have less / The information I've seen is good enough.

Q23: Which of these settings, if any, made you concerned about your privacy? (select all that apply)

Answers: Cross App Data Sharing (BuyNow to RightNow) / Data sharing with 3rd Parties / Data usage for AI Training / None

Q24: Which of these settings, if any, were for something you wanted more information about before making a selection?

Answers: Cross App Data Sharing (BuyNow to RightNow) / Data sharing with 3rd Parties / Data usage for AI Training / None

(if user selected something other than “None” in above question, display below question)

Q24.1: For the question '*Which of these settings, if any, did you want more information about before making a selection?*,' you selected:

[participant response]

What kind of information did you expect to see?

Answers: [free response field]

Q25: Did having an option to partially opt-out of data sharing (e.g., AI Training, third-party advertising) make you feel more in control?

Answers: Strongly Agree / Agree / Neither agree nor disagree (Neutral) / Disagree / Strongly Disagree

Q26: Did you feel that any of the RightNow app settings encouraged you to give up your privacy?

Answers: None of them encouraged me to give up my privacy / Some of them encouraged me to give up my privacy / Most or all of them encouraged me to give up my privacy

[Q26.1 was only shown if the participant selected “Some of them encouraged me to give up my privacy” or “Most or all of them encouraged me to give up my privacy” in response to Q26]

Q26.1: Please tell us what made you feel that the app was encouraging you to give up your privacy.

Answer: [free response field]

Q27: How similar was your approach to making decisions during the setup process to what you would do if installing an app like this in real life?

Answers: In real life I would have spent more time reviewing the options or thinking about which settings to select / In real life I would have taken a similar approach to what I did in this study / In real life I would have spent less time reviewing the options or thinking about which settings to select / Other (Please explain): [free response field]

Q28: Which of the following best describes how you typically select privacy-related settings for social media apps?

Answers: I carefully consider all the choices before making a decision / I try to choose the best choices for me as quickly as I can / I make random choices / Other (Please explain): [free response field]

7) Demographics & Feedback

Please answer the following demographic questions.

Q29: What is your age in years?

Answers: 18-24 / 25-34 / 35-44 / 45-54 / 55-64 / 65 or older / Prefer not to answer

Q30: How do you describe your gender identity?

Answers: Male / Female / Non-binary / Prefer to self describe: [free response field] / Prefer not to answer

Q31: How do you describe your race or ethnic identity? (You may select more than one option.)
[participants can select multiple options]

Answers: White / Black or African American / American Indian or Alaska Native / Asian / Native Hawaiian or Pacific Islander / Hispanic and/or Latino/Latina/Latinx / Prefer to self describe: [free response field] [*exclusive] / Prefer not to answer [*exclusive]

Q32: What is the highest level of education you have received/completed or are working towards?

Answers: Less than high school diploma / High school graduate (high school diploma or equivalent including GED) / Some college but no degree / Associate degree in college / Bachelor's degree / Master's degree / Doctoral degree / Professional degree (JD, MD) / Other, please specify: [free response field] / Prefer not to answer

Q33: Do you have work or educational experience in a computer-related field (e.g. computer science, IT)?

Answers: Yes / No

Q34: Which of the following actions have you taken in the past 3 months while using mobile apps? (select all that apply)

Answers: Declined app permissions (e.g., location, contacts, camera) / Adjusted data sharing settings within apps / Deleted or uninstalled apps due to privacy concerns / Used privacy-enhancing features (e.g., app tracking protection) / Reviewed app privacy policies / Limited ad personalization / None of the above / Other, please specify: [free response field]

Q35: What was your approximate household income in 2023? (based on all household members before tax)

Answers: Less than \$40,000 / \$40,000 to \$59,999 / \$60,000 to \$79,999 / \$80,000 to \$99,999 / \$100,000 to \$149,999 / \$150,000 or more / Prefer not to answer

Q36: If you have any other thoughts or feedback about this survey or the information you viewed, please let us know here (optional)

[free response field]

Thank you for your participation in this survey.
We will send your \$3.00 payment through Prolific.