

Designing Privacy Choice in Generative AI Chatbot Ecosystems

Lanjing Liu*

Department of Computer Science
Johns Hopkins University
Baltimore, Maryland, USA
lliu153@jhu.edu

Allen Yilun Lin

Google Inc.
Mountain view, California, USA
allenylunlin@google.com

Xinran Adeline Li*

Department of Computer Science
Johns Hopkins University
Baltimore, Maryland, USA
xli436@jh.edu

Yaxing Yao

Department of Computer Science
Johns Hopkins University
Baltimore, Maryland, USA
yaxing@jhu.edu

Abstract

Generative AI (GenAI) is evolving from standalone tools to interconnected ecosystems that integrate chatbots, cloud platforms, and third-party services. While this ecosystem model enables personalization and extended services, it also introduces complex information flows and amplifies privacy risks. Existing solutions focus on system-level protections, offering little support for users to make meaningful privacy choices. To address this gap, we conducted two vignette-based survey studies with 486 participants and a follow-up interview study with 16 participants. We also explored users' needs and preferences for privacy choice design across both GenAI personalization and data-sharing. Our results reveal paradoxical patterns: participants sometimes trusted third-party ecosystems more for personalization but perceived greater control in first-party ecosystems when data was shared externally. We discuss design implications for privacy choice interfaces that enhance transparency, control, and trust in GenAI ecosystems.

CCS Concepts

• Security and privacy → Usability in security and privacy.

Keywords

Usable Privacy, Generative AI, User Interface Design

ACM Reference Format:

Lanjing Liu, Xinran Adeline Li, Allen Yilun Lin, and Yaxing Yao. 2026. Designing Privacy Choice in Generative AI Chatbot Ecosystems. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3772318.3790809>

1 Introduction

Generative AI (GenAI) is going through a fundamental shift from standalone tools to operating as ecosystems [47]. Rather than functioning in isolation, GenAI ecosystems are increasingly embedded

in interconnected environments that include chatbots, cloud platforms, and third-party services [12]. These GenAI ecosystems make it possible for personalization and integration with extended services, but they also create complex information flows. Data may move inward, as users disclose personal details (e.g., calendar) to personalize the response, and outward, as the ecosystem shares information with external providers for extended services (e.g., ChatGPT using the Expedia plugin to book a flight). Because of the complex data flows, this ecosystem model, while powerful, amplifies the longstanding privacy challenges and introduces new ones at the same time.

For example, these systems often collect and retain large volumes of personal data, raising risks related to unauthorized data collection, data storage, data leakage, and unintended inference [21, 29]. Meanwhile, GenAI's ability to integrate across platforms blurs boundaries between services, which further creates questions about who can access what information, for how long, and for what purposes [16, 40, 55]. Users also have various levels of privacy concerns about GenAI ecosystems [3, 16, 55]. Yet, existing GenAI ecosystems rarely provide meaningful interfaces for users to make their privacy choices. Existing measures to address the aforementioned privacy issues primarily emphasize model-centric techniques (e.g., data sanitization [23, 34, 40]) to reduce the privacy risks at the system level. Limited prior work has examined how to support users' need to make privacy choices when using the GenAI ecosystem.

Motivated by this critical gap, in this paper, we ask the following research question: *How can privacy choice be designed in the GenAI ecosystem?* The goal is to explore the design space of privacy choice in GenAI ecosystems, as it is unclear whether existing design spaces for privacy notice and choice [14, 39] still apply in the GenAI context. To address this question, we focused on two representative GenAI interaction scenarios, each representing a unique type of information flow: 1) *upstream personalization scenario*, where GenAI requires access to users' data (e.g., calendar) to support personalized services, and 2) *downstream data sharing scenario*, where GenAI shares users' data to external service providers for extended capabilities (e.g., hotel booking). We aim to examine how key design factors (i.e., timing of the privacy choice and types of the GenAI ecosystem) influence users' concern levels, their trust in the GenAI ecosystem, their willingness to use the services, and their perceived sense of control over their data.

*The first two authors contributed equally to this work.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3790809>

We conducted two vignette-based survey studies with 486 participants (Phase 1) and a follow-up interview study with 16 participants (Phase 2). Our findings highlight paradoxical patterns, i.e., participants sometimes trusted third-party ecosystems (e.g., ChatGPT) more for personalization compared to first-party ecosystems (e.g., Gemini), but they perceived greater control in first-party ecosystems than in third-party ecosystems when data was shared to obtain extended services (e.g., booking a flight via Google Flight). We also noticed that the timing of privacy choice (e.g., at-setup vs. just-in-time) had strong effects on users' perceived control and concerns (e.g., participants were more suspicious about just-in-time privacy choice).

This paper makes two key contributions. First, we provide a scoped investigation of users' perceptions, expectations, and concerns regarding privacy choice interfaces in GenAI ecosystems, using the timing of privacy choice as a focused entry point for systematic exploration. We further identified how the types of ecosystem and timing of privacy choice could shape users' sense of privacy control, trust, and willingness to engage. Our findings highlight the paradoxes in users' perceptions. Second, we draw design implications for building usable and trustworthy privacy choice interfaces that empower users in GenAI ecosystems.

2 Related Work

2.1 GenAI Ecosystems

GenAI ecosystems are widely applied in diverse domains, including healthcare [51], finance [4], education [27, 53], and personal counseling [44, 50]. Unlike traditional standalone software, GenAI ecosystems often operate across multiple platforms, combining GenAI capabilities with services such as cloud storage, personalization engines, and third-party APIs. These systems are typically structured as either first-party ecosystems (e.g., Google Gemini, the same company provides both the AI and the platform) or third-party ecosystems (e.g., ChatGPT integrated into external applications), each with distinct data governance models and user expectations [47]. The growing deployment of GenAI ecosystems has heightened the urgency of ensuring Responsible AI practices that prioritize fairness, transparency, accountability, and user empowerment [6, 29]. A central concern in responsible AI is informational fairness, ensuring users understand and meaningfully consent to how their data is used [11, 13]. This is especially relevant in GenAI ecosystems, where personalization and integration are often opaque. Privacy options are an important approach to help users make informed decisions [19, 20]. There remains a gap in understanding how to design privacy options within GenAI ecosystems.

2.2 Privacy Issues of GenAI Ecosystems

As GenAI ecosystems own more capabilities, their growing data requirements intensify existing privacy risks [28], including data collection, processing, dissemination, and invasion risks [29]. Users face a complex data ecosystem. Unlike traditional technologies, GenAI ecosystems can collect data in a proactive way, influence disclosure behaviors through carefully crafted prompts [30], and exploit psychological effects such as nudging or framing [48]. Beyond data collection, GenAI ecosystems' ability to memorize and store user data introduces new risks of data leakage. These models retain

extensive information to enhance personalization and can infer personal attributes such as location, income, and gender [43]. Additionally, collected data is often processed not only within the GenAI ecosystems but also shared with third-party service providers for real-time information access. Instruction-tuned models are particularly vulnerable, as adversaries can manipulate them to bypass privacy-protecting mechanisms [31].

Most privacy risk mitigation strategies in AI technologies focus on model-centric approaches, such as removing private data from training datasets [23, 34] and employing differentially private training techniques [32]. However, current AI product designs frequently fail to prevent users from inadvertently disclosing sensitive information. AI agents offer powerful assistance, making the benefits of data disclosure tangible to users, while the associated privacy risks remain abstract and difficult to assess. This asymmetry can lead users to unknowingly expose themselves to escalating privacy threats over time [55].

Although users may attempt ad-hoc privacy-protective measures when interacting with GenAI ecosystems, they often lack awareness of privacy risks and the impact of their decisions. Tools like ProPILE help identify privacy leaks in LLMs to raise user awareness [26], but practitioners require AI- and product-specific guidance to implement effective privacy protections [28]. Exploring usable privacy design remains essential to enhancing users' awareness and control over their data within GenAI ecosystems [55].

2.3 Privacy Notice and Choice

Originally developed to ensure legal compliance, privacy notices have evolved to promote transparency and support users in understanding how organizations collect, use, and share their data [20, 24]. The design of privacy notices and choices, such as their timing [25] and the way they frame data request purposes [33], significantly affects users' attention to and comprehension of privacy information. Thus, beyond regulatory requirements, effective privacy interfaces need to support users in understanding privacy implications and making informed decisions.

Schaub and Cranor emphasized that effective privacy interfaces require four components, including findability, understandability, usability, and usefulness [7]. Building on this foundation, researchers have explored various approaches to improve the user-friendliness of privacy notices, such as making privacy information and choices more salient and accessible [18, 54]. Schaub et al. introduced a comprehensive design space for privacy notices, which includes dimensions such as timing, delivery channel, modality, and user control mechanisms [39].

Beyond raising privacy awareness through notices, empowering users with effective privacy controls plays a critical role in enabling them to manage their personal data and reduce associated risks. Privacy controls give users agency over what information they share, how it is processed, and with whom it is disclosed [20]. Expanding on existing privacy notice frameworks, Feng et al. introduced additional design dimensions—such as control type, functionality, timing, and communication channel—to improve the usability and effectiveness of privacy choices [15]. Researchers have explored a variety of strategies to support informed decision-making through privacy choice design. These include improving cookie consent

interfaces [18] and presenting privacy-relevant information at decision points in clearer, more actionable formats [25]. Liu et al. developed a personalized privacy assistant to support mobile users in managing app permissions using intelligent, context-aware suggestions [35]. Similarly, Reinhardt et al. introduced a visual interactive privacy policy, enriching it with control options and interactive elements to facilitate user engagement [38].

3 Apparatus

In this study, we aim to explore the key design factors for privacy choices in the GenAI ecosystem. To situate our participants in realistic GenAI ecosystem interactions, we used vignettes (i.e., short scenarios that are systematically generated) paired with visual storyboards. Our goal was to capture two fundamental data flow patterns that occur in GenAI ecosystems: (1) upstream data access, where the ecosystem requests user data to personalize its services, and (2) downstream data sharing, where the ecosystem proposes sending users' information to external service providers to complete extended tasks. Accordingly, we designed two scenarios:

- **Upstream Personalization Scenario**, representing situations where data flows from the user into the GenAI ecosystem. In this scenario, the system requests access to personal information (e.g., calendar, email) to deliver personalized responses.
- **Downstream Data-sharing Scenario**, representing situations where data flows from the GenAI ecosystem outwards to third-party services. In this scenario, the system seeks permission to share user-provided information with an external provider (e.g., a booking service) to enable extended functionality.

We then used storyboards rather than real system interaction to provide a consistent and controlled environment for examining users' privacy perceptions and decision-making process. Each storyboard illustrated the task context, data flow direction, interface elements relevant to the privacy choice, and the consequences of the user's decision. In this section, we detail the processes for generating our vignettes and storyboards.

3.1 Vignette Generation

3.1.1 Upstream Personalization Scenario. In the Upstream Personalization Scenario (hereafter, "personalization scenario"), we focus on users' privacy perceptions, expectations, concerns, and decision-making when a GenAI ecosystem requests access to their data to personalize its responses. We identified two factors: (1) **the type of the GenAI ecosystem**, and (2) **the timing of privacy choice**.

We selected the two factors based on prior work on users' mental models of GenAI ecosystems and long-standing findings in usable privacy research. Prior research shows users hold different trust levels and privacy concerns towards the first-party ecosystems (e.g., Google Gemini) versus third-party ecosystems (e.g., ChatGPT) [47]. The differences further shape data governance models and users' perceived privacy controls. Following the categorization, we used Gemini and ChatGPT as representative examples of first-party and third-party GenAI ecosystems, respectively.

The second factor, timing of the privacy choice, was drawn from the established design space for privacy notice and choice [15, 39]. Although timing is only one dimension in these design spaces, it

is particularly consequential in GenAI contexts. Unlike traditional applications where data flows are explicit and predictable (e.g., permission requests in smartphones), GenAI systems often operate with opaque and dynamic data exchanges, making it difficult for users to know when their data is accessed, processed, and shared. As such, the timing at which a privacy choice is presented to the users can significantly influence users' understandings and perceptions of the data practices and whether they feel empowered to control their data. In the meantime, the broader design space of GenAI privacy disclosure is underexplored. Given the complexity of GenAI ecosystems, including other factors as suggested by prior work (e.g., modality, granularity) [15, 39] would risk introducing too many confounding factors, limiting our ability to interpret the findings. As a result, narrowing our focus to timing provides a tractable and theoretically grounded entry point for systematic investigation. As such, in this study, we focus solely on timing as an entry point for systematic investigation. We selected the following two representative timing points as they captured the critical moments when interacting with GenAI ecosystems (Table 1 summarizes the two factors and their description):

- **At setup:** Users make privacy decisions based on their prior expectations and initial impressions of the system.
- **Just-in-time:** Users make privacy decisions based on both their impression of the system and current interaction.

3.1.2 Downstream Data-sharing Scenario. In the Downstream Data-sharing Scenario (hereafter, "Data-sharing Scenario"), we examined users' privacy-related reactions when a GenAI ecosystem proposes sharing users' data with an external service provider for extended services. We identified three factors: (1) **the type of the GenAI ecosystem**, (2) **the timing of privacy choice**, and (3) **personalization toggle** that may influence users' perspectives and needs. The first two factors mirror those examined in the personalization scenario, allowing us to compare how the type of GenAI ecosystems and timing of privacy choice influence user perceptions across different directions of data flow. The third factor, the personalization toggle, reflects a common practice in GenAI ecosystems, where users may enable or disable personalized responses. We operationalized personalization by varying whether the GenAI ecosystem incorporated users' personal information when generating its responses, as described below.

- **Personalized:** The system generated outputs that explicitly incorporated the user's personal profile or preferences (e.g., tailored recommendations, customized summaries, or content adjusted based on prior user activity).
- **Non-personalized:** The GenAI produced generic, non-user-specific outputs that did not rely on users' personal information or prior behaviors.


Thus, the toggle in each scenario altered the scope of personal information that the GenAI ecosystem could potentially use and expose. By presenting participants with both personalized and non-personalized variants, we examined how personalization influenced users' perceptions of the downstream data sharing and whether they were willing to allow the GenAI ecosystems to share information with external providers. The details of the factors and their levels are presented in Table 2.

Table 1: Factors for personalization scenario varied between vignette scenarios, levels of the factors presented in scenarios, and descriptions of each factor.

Factor	Levels	Description
Type of GenAI Ecosystem	First-party GenAI Ecosystems; Third-party GenAI Ecosystems	The organizational relationship between the GenAI system and the broader platform or service it operates within.
Timing of Privacy Choice	At-setup; Just-in-time	The point during a user's interaction with a GenAI ecosystem when privacy-related notice and choice are presented.


Table 2: Factors for data sharing scenario varied between vignette scenarios, levels of the factors presented in scenarios, and descriptions of each factor.

Factor	Levels	Description
Type of GenAI Ecosystems	First-party GenAI Ecosystems; Third-party GenAI Ecosystems	The organizational relationship between the GenAI system and the broader platform or service it operates within.
Timing of Privacy Choice	At-setup; Just-in-time	The point during a user's interaction with a GenAI ecosystem when privacy-related notice and choice are presented.
Personalization Toggle	Non-Personalized; Personalized	The setting in GenAI ecosystems that enables or disables personalization features, allowing the system to tailor responses or services based on the user's personal data (e.g., calendar, email).

**Figure 1: A-i: Privacy choice mock-up for ChatGPT in Personalization Scenario; A-ii: Privacy choice mock-up for Gemini in Personalization Scenario. B-i: Privacy choice mock-up for ChatGPT in Data-sharing Scenario; B-ii: Privacy choice mock-up for Gemini in Data-sharing Scenario.**

3.2 Generating Vignette-based Storyboards

We selected a flight booking task as the shared context for all vignettes and used Google Gemini and ChatGPT to represent first-party and third-party GenAI ecosystems, following prior work [47].

**Figure 2: A: Example Storyboards for Personalization Scenario; B: Example storyboards for personalized condition in Data-sharing Scenario, where the GenAI provides personalized results; C: Example storyboards for non-personalized condition in Data-sharing Scenario, where the GenAI provides general results.**

Booking tasks naturally involve both necessary information disclosure (e.g., travel dates, destinations) and optional personal information (e.g., preferences, calendar). They also require coordination with external service providers (e.g., Google Flight, Expedia). These characteristics make the flight booking task a suitable and realistic setting for examining how users reason about privacy in both the personalization scenario and the data-sharing scenario.

To present the vignettes, we created storyboards (Figure 2), a technique from prior work on scenario-based evaluation [22]. Each storyboard consists of eight frames that illustrate the task context,




Figure 3: An overview of the study procedure.

technical background, interaction process, mock-ups of privacy choice designs, and resulting consequences.

We designed the storyboards to reflect the visual and interaction language of each GenAI ecosystem. We used a cohesive first-party Google identity for the Gemini storyboards and incorporated multiple third-party service logos to convey its ecosystem role and integration model for the ChatGPT storyboards. For each type, we then designed privacy choice interfaces that were consistent with the visual and interaction language.

To ensure consistency and quality of the vignette and the corresponding storyboards, we first developed detailed textual scripts for every vignette, iterated through multiple rounds of internal team review, obtained feedback from our industry partners who are familiar with GenAI product design, and then produced the finalized visual storyboards. These vignettes, along with the storyboards, are used for the surveys and follow-up interviews.

4 Survey Method

Based on the two vignettes, we conducted two vignette-based survey studies separately, with 246 participants in the Personalization Survey and 240 participants in the Data-sharing Survey, after excluding low-quality responses. We designed two vignette-based surveys to explore users' experiences, preferences, and concerns around privacy design in GenAI ecosystems. We drew on prior user-centered research on privacy choice and control [10, 17], and further developed the items to reflect recent findings on how people understand GenAI ecosystems [47]. The complete survey protocol is available in the Appendix A.1. We describe the study procedures for both the Personalization Survey and the Data-sharing Survey in this section. Our institution's IRB approved the study.

4.1 Survey Protocols

4.1.1 Personalization Survey Protocol. The survey consisted of three parts. First, we assessed participants' prior experiences with AI tools, including products used and typical usage scenarios. Second, participants completed all four vignettes representing the full set of conditions in the Personalization Scenario, in which the GenAI ecosystem requested data for personalization. All storyboards were counterbalanced to avoid order effects. Presenting all conditions to each participant helped control for potential biases arising from differences in participants' prior experiences or backgrounds. After each vignette, participants made privacy decisions, including whether to enable personalization, which accounts to link if personalized, what types of data they considered useful (perceived usefulness), and what types they considered too sensitive (perceived data sensitivity). They also explained their reasoning in open-ended responses. Then, participants rated their perceived sense of privacy control (i.e., their subjective feeling of being able

to manage and influence how their data is used), trust, and usage willingness via Likert scales. Third, we gathered feedback on participants' expectations for privacy control (i.e., the types of privacy settings or controls they would like available) and data information presentation, and their privacy-related concerns by using multiple-choice questions. Demographic and technical background questions concluded the survey.

4.1.2 Data-sharing Survey Protocol. Data-sharing Survey followed a similar structure but focused on scenarios where the GenAI ecosystem proposed sharing user data with external services. Similar to the Personalization Survey, each participant was shown all 8 randomized vignettes, ensuring that every participant experienced the full design space. After each vignette, participants completed Likert-scale items assessing their perceived sense of privacy control, trust, and usage willingness. They then answered multiple-choice questions about their expectations for privacy control and presentation design, followed by demographics and technical background.

4.2 Survey Development and Pilot Study

Before running the formal study, we launched three batches of pilot studies to test our screening mechanism, survey flow, and survey logic for both survey studies. We identified any potential errors in displaying the vignettes in the first pilot survey and determined the number of storyboards each participant needed to review in the second pilot survey.

4.3 Participants Recruitment and Eligibility

We implemented both survey studies using Question Pro and recruited our participants through Prolific in May 2025. We framed our studies as *Designing Interfaces about AI Products* to mitigate self-selection bias from pre-existing privacy beliefs as prior work [2, 46]. We distributed the survey multiple times throughout the week, deliberately releasing each batch at varying times of day and on different days to accommodate participants with diverse schedules. We launched the two surveys simultaneously in each distribution batch, but participants were mutually exclusive across the two studies. This approach ensured independent samples while maintaining consistent timing and survey conditions. We also selected the "even distribution" option on Question Pro to ensure an even distribution of all vignettes.

On Prolific, we set the recruitment criteria to filter participants who are at least 18 years old, located in the US, fluent in English, and have an over 97% task approval rate. All participants have recently used AI products. This is to ensure that, when responding to our survey, our participants could build connections with their prior experiences, further increasing our data. All participants would sign the consent form through Question Pro before taking the survey. After participants completed the survey, we filtered out low-quality responses based on several criteria: unusually fast completion times (e.g., under one minute for the main survey), duplicate submissions (identical responses from different Prolific IDs), and open-ended answers that showed clear signs of copy-pasting (e.g., the same text repeated across all questions). Both studies took an average of 18 minutes to complete. Participants received \$3.00 in the Personalization Survey and \$5.00 in the Data-sharing Survey.

4.4 Data Analysis

We collected quantitative and qualitative data via two surveys.

4.4.1 Quantitative data. We defined three ordinal measures, i.e., privacy control, trust changes on the LLM chatbot, and willingness changes to continue using the LLM chatbot on a 5-point Likert scale. To examine the effects of type, timing, and their interaction, we fitted three Cumulative Link Mixed Models (CLMMs) for our analysis [8]. The model included random intercepts for participants to account for repeated measures ($N = 809$). We corrected p-values with Tukey's Honestly Significant Difference (HSD). To examine users' perspectives, expectations, and needs regarding privacy control, the presentation of privacy-related information, and their concerns, we analyzed responses to binary multiple-choice questions using Generalized Linear Mixed Models (GLMMs) [9]. Similar to CLMMs, the GLMMs model fixed effects of type and timing. Similarly, we adjusted p-values with the Tukey Honestly Significant Difference (HSD) correction to control for multiple comparisons. In cases where we observed marginal patterns ($0.05 < p < 0.10$), we report them only as **exploratory observations** that may suggest potential relationships but do not constitute. We include these patterns primarily to inform directions for future work, as this study serves as an initial exploration of the design space.

4.4.2 Qualitative data. For each open-ended question from the survey studies, two researchers collaboratively coded the responses. They conducted multiple rounds of discussions to reach a full agreement and generated an initial codebook. Then, one researcher coded all data independently using the initial codebook. Upon completion, all authors reviewed the codes together, discussed any disagreements, and refined the codebook as needed until a full agreement was reached. As we collaboratively completed the coding in multiple rounds and reached a complete agreement, the inter-coder agreement was not required [37]. The codebook is available in the supplement materials.

4.5 Ethical Considerations

To mitigate potential privacy concerns, we used storyboards to present the vignettes, rather than collecting data from real-world usage scenarios. We did not collect any personally identifiable information in the study. The only identifier for our participants is their unique IDs from Prolific, which consist of random numbers and letters. Additionally, when we rejected low-quality responses, we provided a short explanation to the participants.

5 Results from Personalization Survey

In this section, we present participants' preferences for privacy choice timing and their expectations for privacy choice when facing different types of GenAI ecosystems in the Personalization Scenario. We first show participants' preferences for the privacy choice timing along with descriptive data. We next describe how GenAI ecosystem types shaped their privacy decisions, information needs, and control expectations.

5.1 Demographic Information

We received 246 valid responses from the Personalization Survey. Our participants represent diverse backgrounds in terms of their

genders, ages, technical levels, etc. The demographic information can be found in Table 3.

5.2 Timing

5.2.1 The timing of privacy choice influences participants' perceived usefulness for personalization. We first asked participants' preferences for different data sources for the GenAI ecosystem personalization. Participants could select multiple choices, including Memory in the Gen chatbot, Google Account, and Microsoft Account. We also did not observe a significant effect on their willingness to integrate a Microsoft Account from the timing of the privacy choice among ChatGPT users. Overall, the timing of the privacy choice did not significantly influence participants' selection of personalization sources. Next, we asked participants which types of data they considered useful and too sensitive for personalization purposes, offering options including Calendar, Email, Contacts, and Photos. Participants viewed "Contacts" as useful for personalization with just-in-time privacy choice ($\chi^2(1, N = 984) = 4.90, p = .029$). Regarding sensitivity, just-in-time privacy choice marginally increased perceptions of Email as sensitive ($\chi^2(1, N = 984) = 3.43, p = .064$), without reaching statistical significance. Notably, the GLMM models could not be reliably fitted for "Calendar" and "Photos" due to insufficient response variance among participants.

5.2.2 At-setup privacy choice increases participants' perceived sense of privacy control. We asked participants about their perceived sense of privacy control provided by the privacy choice. Participants reported a stronger sense of control when the choice was presented at setup rather than just in time. ($\chi^2(1, N = 984) = 4.02, p = .045$).

5.2.3 The timing of privacy choice did not impact participants' perceived trust in the GenAI ecosystems. We asked participants to rate their perceived changes in their trust in the ability of GenAI to handle their data responsibly. The CLMM model showed that the timing of privacy choice didn't have a significant impact on participants' trust ($\chi^2(1, N = 984) = 0.06, p = .800$). Most participants reported "No change" in trust (62.60%, $n = 616$) while 24.08% participants ($n = 237$) selected "I trust it somewhat more".

5.2.4 The timing of privacy choice did not impact participants' willingness to use GenAI ecosystems. We did not observe a significant effect on participants' willingness from the timing of privacy choice ($\chi^2(1, N = 984) = 0.53, p = .466$). Most participants reported "No change" in their willingness to continue using the system (59.04%, $n = 581$), indicating relatively stable intentions across conditions.

5.2.5 The timing of privacy choice did not impact participants' need for privacy control. We examined participants' preferred privacy controls when GenAI ecosystems requested access to their data. The GLMM analyses did not reveal a significant effect of timing on participants' control preferences. This suggests that participants' desired privacy controls may be generally consistent regardless of when the privacy choice is presented.

The three most frequently selected control options were: "Choose in advance which types of data ChatGPT can or cannot access (e.g., block location sharing)" (70.63%, $n = 695$), "A dashboard to track who received my data" (55.18%, $n = 542$), and "Real-time approval before each data transfer" (54.57%, $n = 537$). The high selection rate for

Table 3: Participants’ demographic information for personalization survey

Gender		Age		Education		Employment		Tech Level	
Female	46.3%	18-24 years old	4.5%	High school or below	34.1%	Working full-time	64.2%	Ultra Nerd	8.1%
Male	51.2%	25-34 years old	26.8%	Bachelor degree	49.2%	Working part-time	16.7%	Technically Savvy	60.6%
Non-binary	1.2%	35-44 years old	25.6%	Master degree	10.6%	Unemployed and looking for work	4.1%	Average User	30.1%
Prefer not to answer	1.2%	45-54 years old	24.0%	Doctoral degree	2.8%	A homemaker or stay-at-home parent	4.5%	Luddite	1.2%
		55+ years old	17.9%	Prefer not to answer	3.2%	Student	2.8%		
		Prefer not to answer	1.2%			Retired	5.3%		
						Other	2.4%		

pre-authorization and real-time approval of data types indicates that participants value granular and immediate control over their personal information. This suggests a desire for GenAI ecosystems that allow users to proactively define privacy boundaries, rather than relying solely on passive or after-the-fact controls.

5.2.6 Just-in-time privacy choice increased participants’ desire for some privacy-related information. We then investigated participants’ needs for privacy information through a multiple-choice question. When the privacy interface was shown just in time, participants showed a greater need to control the “specific types of data being accessed” ($\chi^2(1, N = 984) = 7.75, p < .001$). We also observed a marginal trend towards a greater need to control “how long the data will be stored” ($\chi^2(1, N = 984) = 2.57, p = .077$).

Participants prioritized transparency around data sharing, access rights, and specific data types collected. “Whether data will be shared with third parties (e.g., advertisers, partner companies)” is the top option (59.76%, $n = 588$). “Who can access the data (e.g., only ChatGPT, OpenAI, third-party providers)” is the second (56.40%, $n = 555$), and “Specific types of data being accessed (e.g., emails, calendar entries, files in Drive)” is the third (54.98%, $n = 541$).

5.2.7 At-setup privacy choice reduced participants’ privacy concerns. Finally, we asked participants to report their privacy concerns when the GenAI ecosystem requests access to their data for personalization. When privacy choice presenting at setup, participants were significantly less likely to report concerns about *unauthorized data collection* ($\chi^2(1, N = 984) = 6.61, p = .010$), *potential misuse by the AI* ($\chi^2(1, N = 984) = 8.37, p = .004$), *loss of control over shared data* ($\chi^2(1, N = 984) = 8.42, p = .004$), and *inability to track how data is used* ($\chi^2(1, N = 984) = 18.85, p < .001$). We also observed marginal reductions in concerns about *the lack of transparency about what data is collected* ($\chi^2(1, N = 984) = 3.72, p = .054$) and data sharing with untrusted third parties ($\chi^2(1, N = 984) = 3.00, p = .083$), though these effects did not reach conventional statistical significance.

Among the concerns when the GenAI ecosystem requests access to their data for personalization, participants were mostly concerned with “data shared with untrusted third parties” (71.85%, $n = 707$), “unauthorized data collection” (57.31%, $n = 564$), and “lack of transparency about what’s collected” (54.37%, $n = 535$).

5.3 Type of GenAI Ecosystems

5.3.1 Participants preferred to use chatbots’ memory for personalization in third-party ecosystems, whereas they preferred to use data

from the same company for personalization in first-party ecosystems. Participants showed a stronger preference for Memory ($\chi^2(1, N = 984) = 6.46, p = .011$) when facing ChatGPT than Gemini. While those interacting with Gemini favored connecting their Google Account over ChatGPT ($\chi^2(1, N = 984) = 19.48, p < .001$). Regarding sensitivity, Gemini decreased the number of participants who considered “Email” sensitive ($\chi^2(1, N = 984) = 5.95, p = .015$). We didn’t find a significant effect on participants’ perceived usefulness of data resources from the type of GenAI ecosystem.

5.3.2 Third-party ecosystems increased participants’ perceived sense of privacy control. Regarding perceived sense of privacy control, participants perceived ChatGPT providing more sense of privacy control compared to Gemini ($\chi^2(1, N = 984) = 3.83, p = .050$).

5.3.3 Type of GenAI ecosystem did not affect participants’ perceived trust and their willingness to use the GenAI ecosystem. The CLMM model showed that the type of GenAI ecosystem didn’t have a significant impact on participants’ trust ($\chi^2(1, N = 984) = 0.02, p = .886$) as well as participants’ willingness from the type of GenAI ($\chi^2(1, N = 984) = 0.26, p = .610$).

5.3.4 Type of GenAI ecosystem did not affect participants’ need for privacy control. Among the options for privacy concerns, we did not observe significant effects of either type on participants’ control preferences from the type of GenAI ecosystem. This suggests that users’ expectations for managing their data remain relatively stable regardless of whether they interact with a first-party or third-party GenAI ecosystem.

5.3.5 Third-party ecosystems increased participants’ need for some privacy-related information. We found that when using ChatGPT, participants were more likely to request information related to *the purpose of data use* ($\chi^2(1, N = 984) = 12.99, p < .001$), *who can access the data* ($\chi^2(1, N = 984) = 10.53, p = .001$), and *potential risks of exposure* ($\chi^2(1, N = 984) = 4.66, p = .031$). This suggests that third-party ecosystems like ChatGPT trigger stronger privacy concerns, prompting users to seek more granular control over how their data is used, who can access it, and the potential risks associated with sharing.

5.3.6 First-party ecosystems and participants’ privacy concerns. The GLMM model indicated that Gemini has marginally decreased participants’ concern about *the loss of control over shared data* ($\chi^2(1, N = 984) = 3.07, p = .080$), though this effect was not statistically significant. While this effect did not reach the level of significance, it

suggests a possible sensitivity to just-in-time privacy choice. The results indicate that just-in-time privacy choice increased participants' awareness and concern about various privacy risks compared to providing them at setup.

6 Results from Data-sharing Survey

In this section, we report participants' preferences for privacy choice timing and expectations when facing different GenAI ecosystem types and personalization toggles in the Data-sharing Scenario. We first present participants' preferences for the privacy choice timing along with descriptive data. Then, we describe how GenAI ecosystem types and personalization toggles shaped their privacy decisions, information needs, and control expectations.

6.1 Demographic Information

We received 240 valid responses for the Data-sharing Survey. The participant demographic composition in the Data-sharing Survey closely mirrors that of the Personalization Survey in terms of gender, age distribution, technical background, and AI usage. Table 4 summarizes the demographic characteristics across both studies.

6.2 Timing

6.2.1 At-setup privacy choice and participants' perceived sense of privacy control. Similarly, in each vignette, we asked participants about their perceived sense of privacy control of the interface in the scenario. We observed a marginally significant effect of timing ($\chi^2(1, N = 1920) = 3.46, p = .063$), indicating a trend towards a higher perceived sense of control with at-setup choices compared to just-in-time ones.

6.2.2 Just-in-time privacy choice and participants' perceived trust. We also asked participants to evaluate how their trust in the GenAI ecosystem's ability to handle their data responsibly changed as a result of the interface they just interacted with in the Data-sharing Survey. The CLMM model showed that just-in-time privacy choices marginally increase participants' perceived trust change ($\chi^2(1, N = 1920) = 3.12, p = .078$), suggesting a possible, but not statistically significant, increase in trust with timely, context-specific prompts. The majority of participants selected "No change" (51.72%, $n = 993$). Following that, 22.14% participants ($n = 425$) selected "I trust it somewhat more", and 17.50% participants ($n = 336$) selected "I trust it somewhat less".

6.2.3 The timing of privacy choice did not impact participants' willingness to use GenAI ecosystems. We then asked participants how the interface changed their willingness to continue using the GenAI ecosystem. We did not find a significant effect from the timing of the privacy interface ($\chi^2(1, N = 1920) = 2.57, p = .110$). Most responses are "No change" (51.51%, $n = 989$) in their willingness to continue using the GenAI ecosystem. The second selected response was "I'm somewhat more willing to use it" (24.21%, $n = 465$).

6.2.4 Just-in-time privacy choice increased participants' need for privacy control. We asked participants about their need for privacy control when the GenAI ecosystem shares information with external service providers through a multiple-choice question. We observed that participants show more interest in "real-time approval before each data transfer" when the interface that provides privacy

choice is present just in time ($\chi^2(1, N = 1920) = 7.69, p = .006$). This indicates that just-in-time interfaces may prompt users to seek more granular, task-specific control over their data. For the privacy control needs, consistent with Personalization Survey (Section 5.2.5), the three most selected options were "choose in advance which types of data ChatGPT can or cannot access" (66.82%, $n = 1283$), "real-time approval before each data transfer" (47.66%, $n = 915$), and "a dashboard to track who received my data" (47.29%, $n = 908$).

6.2.5 At-setup privacy choice and participants' need for privacy-related information. Regarding participants' need for privacy information presentation, we observed that at-setup privacy choices were associated with a marginally higher need for information about "potential risks of data exposure" ($\chi^2(1, N = 1920) = 3.45, p = .063$) when facing with at-setup privacy choice than just-in-time ones. While this effect did not reach conventional significance levels, it suggests that just-in-time privacy choice may provide participants with sufficient contextual information, potentially reducing their perceived need for additional risk-related details. It is noted that for the option "physical location of data storage (e.g., servers in specific countries)", the model showed convergence issues, likely due to near multicollinearity among predictors or variables on different scales, which can affect the stability of parameter estimation. When the GenAI ecosystem notifies participants about connecting to third-party services, the most commonly selected types of presented information were "who can access the data (52.97%, $n = 1017$)", "whether data will be shared with third parties (52.24%, $n = 1003$)", "specific types of data being accessed (51.82%, $n = 995$)".

6.2.6 At-setup privacy choice and participants' privacy concerns. The GLMM model for privacy concerns indicated a marginal trend toward decreased concern over the lack of transparency about what data is collected when privacy choice appeared at setup ($\chi^2(1, N = 1920) = 3.01, p = .083$). This pattern suggests a possible, though not statistically significant, effect whereby presenting privacy choice at setup may help users anticipate data use and feel more informed about how their information is handled. During the process in which the GenAI ecosystem notifies participants about connecting to third-party services, the most frequently reported concerns were "data shared with untrusted third parties" (62.92%, $n = 1208$), "unauthorized data collection" (55.57%, $n = 1067$), and "lack of transparency about what's collected" (53.13%, $n = 1020$).

6.3 Type of GenAI Ecosystem

6.3.1 First-party ecosystems and participants' sense of privacy control. Our CLMM model suggests a marginally significant effect of type on participants' perceived sense of privacy control ($\chi^2(1, N = 1920) = 3.27, p = .066$), with participants reporting that Gemini gives them slightly higher control than ChatGPT. Although not statistically significant, this likely reflects that first-party ecosystems like Gemini feel more integrated and predictable, giving slightly higher perceived control than third-party systems like ChatGPT.

6.3.2 Type of GenAI ecosystem did not impact participants' perceived trust and privacy concerns. Regarding the perceived trust in GenAI ecosystems, we did not observe a significant effect on participants' trust from GenAI ecosystem type ($\chi^2(1, N = 1920) = 2.42, p = .120$), suggesting that participants' trust in GenAI ecosystems

Table 4: Participants' demographic information for data-sharing survey

Gender		Age		Education		Employment		Tech Level	
Female	44.2%	18–24 years old	3.3%	High school or below	38.3%	Working full-time	61.7%	Ultra Nerd	7.5%
Male	55.8%	25–34 years old	27.1%	Bachelor degree	44.2%	Working part-time	17.1%	Technically Savvy	64.6%
		35–44 years old	24.2%	Master degree	12.9%	Unemployed and looking for work	8.3%	Average User	27.5%
		45–54 years old	24.6%	Doctoral degree	2.5%	A homemaker or stay-at-home parent	2.9%	Luddite	0.4%
		55+ years old	20.8%	Prefer not to answer	2.1%	Student	1.2%		
						Retired	5.0%		
						Other	3.8%		

may depend less on whether it is a first-party or third-party system. Similarly, we observed no significant effect of ecosystem type on privacy concerns, indicating that ecosystem type alone does not strongly influence users' control preferences or worries about data handling.

6.3.3 First-party ecosystems increased participants' willingness to continue using GenAI ecosystems. The CLMM revealed a significant main effect of type of GenAI ecosystem ($\chi^2(1, N = 1920) = 4.97$, $p = .026$), with participants reporting higher willingness under Gemini compared to ChatGPT. The difference may reflect participants' greater familiarity or trust with Gemini as a first-party ecosystem, rather than the GenAI itself.

6.3.4 Third-party ecosystems increased participants' need for privacy control. We also observed that ChatGPT has significantly increased participants' needs about "automatic expiration of shared data" ($\chi^2(1, N = 1920) = 5.64$, $p = .018$). It indicated that participants were less confident in how the third-party ecosystem would manage or retain their information over time.

6.3.5 Participants' mixed need for privacy-related information. We observed that participants expressed greater needs for the information about "purpose of using the data" when using Gemini compared to using ChatGPT ($\chi^2(1, N = 1920) = 5.17$, $p = .023$). They also expressed increased interest in information about "security measures to protect the data" ($\chi^2(1, N = 1920) = 6.66$, $p = .010$) when using ChatGPT. The model of option "physical location of data storage (e.g., servers in specific countries)" also showed convergence issues for the type of GenAI ecosystems.

6.4 Personalization Toggle

6.4.1 Personalization toggle did not significantly influence participants' sense of privacy control, perceived trust, willingness to use the GenAI ecosystem, or privacy-related control preferences. We did not find a significant effect of the personalization toggle on perceived privacy control ($\chi^2(1, N = 1920) = 2.66$, $p = .103$, participants' trust ($\chi^2(1, N = 1921) = 1.83$, $p = .176$), or changes in willingness to use the GenAI ecosystem ($\chi^2(1, N = 1920) = 1.28$, $p = .258$). Similarly, among the options for privacy control, the personalization toggle did not produce any significant effects.

6.4.2 Personalization toggle increased participants' need for privacy-related information. In terms of privacy-related information presentation, the GLMM model revealed a significant effect of personalization toggle on participants' interest in the "purpose of using the data" ($\chi^2(1, N = 1920) = 4.64$, $p = .031$), suggesting that participants have a greater need to understand why their data is being used when interacting with a personalized GenAI ecosystem.

6.4.3 Participants' mixed privacy concerns. When the GenAI ecosystem was personalized, participants expressed greater concern about "inability to track how data is used" ($\chi^2(1, N = 1920) = 7.74$, $p = .005$), suggesting that personalization increases attention to transparency and data traceability. Surprisingly, participants showed reduced concern about "loss of control over shared data" ($\chi^2(1, N = 1920) = 5.62$, $p = .018$) with personalized GenAI ecosystem.

6.5 Key takeaway from survey results

The survey studies reveal a directional paradox in participants' privacy perceptions and preferences (Table 5). In the Personalization Survey, participants reported less perceived sense of privacy control when entrusting personal data to first-party GenAI ecosystems such as Gemini (Section 5.3), and they expressed a preference for third-party ecosystems like ChatGPT. However, the third-party ecosystems also raised participants' needs for several privacy information presentations, such as who can access the data, and potential risks of exposure. Conversely, in the Data-sharing Survey (Section 6.3), first-party ecosystems (e.g., from Gemini to Google Flights) gave participants a greater perceived sense of privacy control compared to third-party ecosystems. The survey studies also showed a consistent preference for the privacy choice presented at setup. In both survey studies, participants associated the at-setup privacy choice with a stronger perceived sense of privacy control, and the just-in-time privacy choice brought more privacy concerns. Specifically, in the Personalization Survey, just-in-time privacy choice significantly increased all types of participant concerns; in the Data-sharing Survey, they specifically heightened concerns about the lack of transparency regarding data collection.

Taken together, these contradictory patterns suggest that participants hold more nuanced and complex privacy perceptions, concerns, and preferences than what the survey study alone can reveal. To gain deeper insight into the reasoning behind these responses, we conducted a follow-up interview study.

Table 5: Comparison of key privacy constructs across upstream and downstream scenarios.

Construct	Upstream Personalization (User → GenAI)	Downstream Data Sharing (GenAI → External Service)	Notable Cross-Scenario Contradictions
Perceived Privacy Control	Higher for At-setup. Higher for third-party ecosystems.	Marginally higher for first-party ecosystems	Opposite effect of ecosystem type
Trust	No significant effects	No significant effects	Trust stable
Willingness to Use	No significant effects	Higher for first-party ecosystems	Ecosystem type matters only downstream
Control Needs	No significant effects	Just-in-time increases the desire for real-time approval. Third-party ecosystem increases automatic-expiration control	Timing and ecosystem type matter only downstream
Information Needs	Just-in-time increases the need for specific data-type details. Third-party increases the need for purpose and access	Ecosystem-type effects are mixed. Personalization increases the need for purpose.	Purpose-of-use importance shifts across scenarios
Privacy Concerns	At-setup reduces most concerns.	Effects narrower (primarily transparency-related)	At-setup timing reduces concerns upstream but not downstream

7 Follow-up Interview Study

Motivated by these paradoxical patterns, we conducted a follow-up interview study with 16 participants. We presented all vignettes to each participant to support a holistic comparison across the two scenarios. This interview study allowed us to gather deeper, interactive insights that help understand the survey results and reveal the perceptions, concerns, and preferences underlying participants' privacy preferences. In this section, we describe the interview methods and key findings.

7.1 Method

7.1.1 Protocol. The interview consists of three parts. In the first part, we collected participants' demographic information and explored their daily experiences with GenAI chatbots. In the second part, we presented Personalization Scenario storyboards, asking participants to indicate their personalization choice and explain their reasoning. We then asked about their perceived sense of privacy control, trust changes, concerns, and needs related to privacy management. After discussing all storyboards, we asked participants to rank those conditions presented in the storyboards and give reasons. The third part followed the same procedure using Data-sharing Scenario storyboards, probing privacy perceptions, trust, concerns, and needs, followed by condition ranking. All storyboards were counterbalanced to avoid order effects.

7.1.2 Participants Recruitment. Similarly, we recruited our participants through Prolific by framing our studies as "Design Interfaces about AI Products Interview" in July and August 2025. We also distributed the survey multiple times throughout the week. On Prolific, we set the same recruitment criteria as survey studies. We first use the screening survey to filter eligible respondents with prior Gen AI experience. The average completion time of the screening survey was 1 minute 21 seconds, and participants received \$0.2 after

completing the screening survey. After the screening survey, all eligible participants continued to take the interview after signing another consent form. The duration of the interviews ranged from 30 to 60 minutes. Upon interview completion, they would receive another \$25 bonus, making the total compensation \$25.2.

7.1.3 Data Analysis. For the interview data, two co-authors read the interview transcripts several times. Then they coded data from two participants collaboratively and generated two initial codebooks (one for the Personalization Scenario and one for the Data-sharing Scenario). Using the initial codebooks, the two co-authors coded the rest of the data individually. They constantly compared and discussed their codes, resolving any disagreements as they coded, and then updated the codebooks as needed. Upon completion of the coding, the two co-authors cross-checked each other's coding again to ensure full agreement. Then, the co-authors examined and discussed the codebooks and grouped the codes into higher-level themes. The final codebook contains 56 unique codes in 8 themes and 44 unique codes in 7 themes, respectively.

7.2 Demographic Information

We recruited 16 participants for the interview study, representing a diverse range of backgrounds and experiences with both first-party and third-party GenAI ecosystems. All participants were English-speaking and based in the United States. We compared demographics, including technical skills and GenAI experience, and found that interview and survey participants had similar profiles, supporting the relevance of qualitative insights to the quantitative results. We observed data saturation by the 13th interview, but we continued recruitment until 16 participants to ensure saturation. Table 6 summarizes the participants' demographic information.

Table 6: Participants' demographics information for interview study

ID	Age	Gender	Ethnicity	Profession	Technology Skills	GenAI Chatbots Usage	Frequency Usage
P01	55	Male	White	Retired	Above average	ChatGPT, Gemini	Once a day
P02	33	Male	Middle Eastern	Q&A tester	Expert	Deepseek, ChatGPT, Gemini	Several times a day
P03	33	Male	White	Construction worker	Pretty well	Deepseek, ChatGPT, Gemini, DALLÉ	Several times a day
P04	43	Male	White	Commerce business owner	Proficient	ChatGPT, Gemini	Once a day
P05	54	Female	White	College professor	Pretty well	ChatGPT, Gemini	Several times a day
P06	54	Male	Black	Teacher	Pretty well	ChatGPT, Gemini	Several times a week
P07	32	Male	White	SVP assistant	Pretty well	ChatGPT, Gemini	A few times a week
P08	24	Female	European/Asian	Semiconductor business owner	Pretty well	ChatGPT, Gemini, Deepseek, DeepL	Several hours a day
P09	35	Male	White	Retail sales	Above average	ChatGPT, Gemini	Several times a day
P10	30	Female	Asian	Social worker	Pretty well	ChatGPT, Gemini, Snapchat AI	A few times a week
P11	32	Male	Asian	Data engineer	Proficient	ChatGPT, Gemini, Meta AI	Several times a day
P12	45	Female	White	AI trainer	Above average	ChatGPT, Gemini, Copilot, Preplexity	Several times a day
P13	41	Male	White	College professor	Above average	ChatGPT, Gemini, Copilot	A few times a week
P14	31	Non-binary	White	Teacher	Average	ChatGPT, Gemini, Copilot	A few times a week
P15	30	Male	White/Asian	Digital marketing	Proficient	ChatGPT, Gemini, Jasper AI	Several times a day
P16	55	Female	White	Full-time mom	Pretty savvy	ChatGPT, Gemini, Copilot	Several times a day

7.3 Personalization Scenario: Information input to GenAI chatbot for personalization

7.3.1 Timing. At-setup privacy choice provided general privacy control in advance and supported a natural interaction, but they could be intrusive.

Overall, participants viewed the at-setup privacy choice as general control and management in advance. P02 remarked, “If you get the pop-up before you do the search, basically this is for the full experience, and not only your question...After you ask about the itinerary, you ask about it. It’s still going to look through all of your data.” P07 extended this metaphor, describing the at-setup privacy choice as a gate before interacting with the GenAI chatbot: “Especially because you do it before anything gets done. Again, it’s like a metaphorical gate of saying by crossing this gate, this is what’s gonna happen.” Several emphasized the at-setup privacy choice because these options clarified how data would be used before any decisions. As P08 explained, “Prompt asking me if I would like a more personalized approach. Because I wouldn’t waste time inputting itinerary dates and whatnot.” In addition, participants valued the at-setup privacy choice because it required less effort than just-in-time options when managing their privacy. As P08 explained, “Because it’s a lot more convenient to give it access to what it needs before I prompt it... whereas if I just allow it to have access to my prior information. I would have to... I don’t have to provide as many details.”

Some participants also framed the at-setup privacy choice as a more natural part of the interaction, aligning with their sense of an ongoing conversation with the GenAI chatbot. P01 compared it to the beginning of an interview: “It’s just like this interview, for example, [the privacy choice like] when you started with the introduction, the purpose for being there, and then you ask for consent to record the meeting. It just makes sense to me that you would get the permissions. And consents done in the beginning, so there’s no misunderstanding.” However, some participants perceived at-setup prompts as intrusive. P12 noted, “I might find the prompt before I give it anything. A little bit more intrusive than if they ask me afterwards.” P04 echoed, “I feel like if it pops up before, that might

be something a little scared of... the steps that you want to get to, and then the pop-up appears, so it might not kind of scare people off.”

Just-in-time privacy choice offered task-specific, contextual control, but users perceived it as delayed disclosure. However, interestingly, some participants valued just-in-time privacy choice more because just-in-time provided more specific and relevant control. P02 explained, “It means it’s only relevant to your search. This is something that I’m going to do... And it understands what you’re searching for, and it gives you prompts related to your search.” P16 echoed this point, describing just-in-time prompts as more task-oriented: “Right up front... I want to go to Colorado. But it doesn’t know anything else, necessarily. So I like [just-in-time] better... Because then I can decide. If I wanted to pull my information and do its AI thing, and give me [results].” Additionally, just-in-time privacy choice could avoid scaring people, since they come after they have already completed their initial steps. P04 said, “After the needs, you already got past the steps that you want to get to, and then the pop-up appears, so it might not kind of scare people off.”

Meanwhile, some participants felt that just-in-time privacy choice reduced control once some data had already been shared. As P14 explained, “You already gave it information and data. And then it’s asking, so it could have taken stuff from that initial prompt. You didn’t have the choice.” Just-in-time prompts also triggered suspicion when systems requested additional information after participants had already provided input. P08 noted, “If I already fed it information. And it still needs more information; I would find that a bit suspicious.”

7.3.2 GenAI Ecosystem Types. First-party ecosystems were perceived as richer, more integrated systems, but suspicious regarding their hidden data transfers and uses. During the interviews, participants frequently discussed the information flow within first-party GenAI ecosystems. Some participants preferred first-party GenAI ecosystems because they had longer and more frequent exposure to them, which built trust. P11 explained, “I trust Gemini and Google a little bit more. I have a Google One subscription and cloud subscription. So I just use Google a lot, and I’ve been using

that for many years.” This perspective aligns with the Personalization Survey findings, where participants expressed weaker concern about losing control over shared data when using Gemini compared to ChatGPT (Section 5.3).

Many assumed that first-party systems engage in extensive data transfer and generally hold much information about their users. Participants thought that first-party GenAI ecosystems have more information about them. As P07 explained, “*Since Google’s more silent and has more information about me, I want to be able to have that information pulled out before I do anything.*” Participants also assumed that first-party ecosystems are already connected to other services, which influenced their perceptions of data access. P08 said, “*I like the way that ChatGPT. But Gemini, I find it suspicious that it’s asking for a personalized experience, because it’s already located on Google, and so it already has access to my Google information.*” Participants also believed that once first-party GenAI ecosystems obtained their information, the systems would further develop their user profiles across the ecosystem. P07 explained, “*Because there’s so much people use Google for, like me, using email search. There are probably people who use Google Travel. Outside of using AI, that’s maybe concerning, because it develops more of a profile, versus ChatGPT, where I think it’s just one thing.*”

Thus, participants perceived privacy consent requests for personal information for personalization as over-asking and sometimes confusing when using the first-party GenAI ecosystem. P08 explained, “*because it’s a Google service, it has access to my Google information, but if it’s asking for more Google information. I don’t know what that extra Google information is.*” Participants expressed a preference for a guarantee that their information would remain self-contained within the GenAI ecosystem rather than being shared across other systems of the same company. P04 noted, “*It is similar to what I would do with ChatGPT, the memory in Gemini. Again, having it right in the program and self-contained, assuming it’s not shared with Google. That, since you’re not connecting or installing the account, that’s probably what I’d go with first option [memory in Gemini].*” These perspectives align with the Personalization Survey results (Section 5.2), where participants reported a greater sense of privacy control when using ChatGPT compared to Gemini.

7.3.3 Expectations and Needs: Enhance the Information Transparency and Data Access Control. Participants also expressed specific privacy needs with the GenAI ecosystems, largely aligning with findings from the Personalization Survey. Regarding presentation, participants in the interviews emphasized the importance of clarifying the data access scope (e.g., P10), who the data would be shared with (e.g., P11, P12), and how the data would be used (e.g., P11, P12). These priorities are closely aligned with survey results (Section 7.4.3). Regarding control, participants stressed the ability to manage and restrict data access, which also aligned with the top survey priority (Section 5.2).

7.4 Data-sharing Scenario: Information output from GenAI chatbot to external providers

7.4.1 Timing. At-setup privacy notice provided advanced clarification, but also vague scope and intrusive questioning.

Participants generally viewed the privacy choice presented at setup as providing strong control before any data exchange. As P15

noted, “*I’d rather have that start immediately, be on notice immediately, as I’m starting the whole process. So you’re still willing to maybe manage and control it before you start anything else.*” Similar to earlier findings (Section 7.3.1), participants valued at-setup options for offering clarity and reducing surprise. P08 said, “*Because just having it beforehand. I think that’s a lot more useful and less surprising and more useful.*” By contrast, some participants perceived just-in-time prompts as a delayed disclosure that reduced trust. P14 explained, “*I think if it... times you afterwards, I have less trust in it, because it already took some of my information. Before letting me know how I could be used or giving me a choice.*” However, some participants described the at-setup privacy choice as confusing or intrusive, especially when they appeared disconnected from the task. As P13 remarked, “*It just jumps out to me as a little bit suspicious. It’s not asking me for information about myself. It’s connecting to something externally. Why is it doing that when I haven’t asked it anything or given it any information?*”

Just-in-time privacy notice offered contextual, task-specific control but could disrupt the natural flow of interaction. Consistent with Section 7.3.1, some participants valued just-in-time privacy choice for offering task-specific and contextual control. P02 noted, “*I think in this scenario, definitely option B is going to be clearer. Because it pops up after the prompt, it’s only related to your flight information. This scenario makes a lot of sense.*” Some participants also described just-in-time privacy choice as random and interruptive, disrupting the interaction flow. As P08 explained, “*It [just-in-time] just seems an ill-conceived time, because it doesn’t make sense to have it. After I’ve already prompted it to do something. It would just appear in front of where my generation, or where my prompts’ answer would have been. So I find A [just-in-time] inconvenient, and B [at-setup] a lot more convenient.*”

7.4.2 GenAI Ecosystem Types. First-party ecosystems hide data transfers and uses for profit. Same as Section 7.3.2, some participants assumed that first-party GenAI ecosystems were already interconnected and would use data from the chatbot across other services. As P08 remarked, “*what Google is telling me is something that is already sort of intuitively known.*” Additionally, P07 preferred ChatGPT over Gemini, perceiving less risk of data reuse: “*Probably getting ChatGPT, because I think there’s less risk of my personal data being used, especially because, again, Google’s such a giant company, and my data’s gonna be used for certain services.*” Participants also questioned the motivations behind sharing information with external service providers, suspecting that first-party ecosystems benefited financially from such integrations. P16 captured this skepticism: “*I always find it a little suspicious when Google’s recommending Google Flights, like, are they making a profit out of this?*”

Third-party ecosystems introduced more unknowns and increased the need for transparency. However, participants felt that when using third-party GenAI ecosystems, there was more they needed to know. They wanted clearer disclosures about which external parties received their data and what types of data were being shared. The lack of such information reduced trust, as P11 explained, “*But if they don’t mention the third party, like ChatGPT is not mentioning the third parties. Then it reduces my trust a little. Because I don’t know what they are sharing my data with.*”

7.4.3 Personalization Toggles. Personalization toggles enhanced the relevance and reliability of results. In the interviews, most participants preferred personalization scenarios, as they believed personalization could provide more relevant and reliable outputs. Many participants viewed non-personalized interactions as producing less reliable results. At the same time, participants described being more cautious when personalization was involved, particularly regarding data sharing. For instance, P11 explained that he would choose personalization with Gemini, but for ChatGPT, he would only proceed if third-party partners were explicitly named. P13 further expressed that personalization prompted further questions, such as “*How much of my information is ChatGPT giving to those third parties? And data usage, data storage?*”

7.4.4 Expectations and Needs: transparency of service providers and control of data sharing among different stakeholders. Participants’ privacy expectations and needs largely aligned with the survey findings (Section 6.2). For privacy-related information presentation, participants highlighted the importance of understanding which service providers could access their data and how it would be shared. For privacy control, interview participants emphasized managing data sharing based on specific tasks and controlling which services and service providers could access their data.

7.5 Comparison of Personalization and Data-sharing Scenario

Since we presented all storyboards to participants during the interviews, we observed an interesting phenomenon: some participants shifted their preferences depending on the scenario. Their attitudes toward the timing of privacy choice varied (e.g., P04), and they weighed the importance of timing differently when compared to the type of GenAI ecosystem (e.g., P06, P08, and P14). To better understand these trade-offs, we asked these participants to directly compare the two scenarios.

7.5.1 Perceived boundary between Personalization Scenario and Data-sharing Scenario. Regarding the Personalization Scenario (upstream data access to the GenAI chatbot for personalization), participants viewed the interaction as remaining within the same ecosystem. In contrast, for the Data-sharing Scenario (GenAI chatbot sending users’ information to external service providers), participants emphasized the **change of environment when data moved outside the original system**. They felt that this transition warranted an immediate notice or consent prompt, since it signaled leaving the familiar ecosystem. As P04 explained, “*When it’s about to change connection, to go into a third party, it’s a completely different website. I think it’s just very important that right away you want to show that pop-up. And give people a heads-up or an option. Because you’re navigating away from the website, but where they are, and going somewhere completely different.*” As a result, they considered at-setup prompts less critical. For example, P14 said, “*To show that notice or pop-up right away, just because you’re staying on the same websites, the same ecosystem. You’re just advocating for a different aspect, something that’s non-personalization from the personalization.*”

7.5.2 Perceived Differences in Information Exposure. Participants also distinguished the **information exposure** across the two scenarios. For the Data-sharing Scenario, where data is downstreamed, they appreciated that no additional personal information needed to be entered into the chatbot: “*It’s because it’s not pulling any information from me, so I don’t need to do extra work in order to have it complete a task.*”

8 Discussion

8.1 Unpacking Participants’ (Shifted) Perceptions of Just-in-Time Privacy Choice in GenAI Ecosystems

Interestingly, our results also indicate some inconsistency when compared to the literature. Prior work has suggested that just-in-time notices can help users better understand what information is being collected and when [14, 39]. Some of our participants agreed. However, across both survey and interview studies, more participants preferred the at-setup privacy choice. Our interview participants explained that at-setup notices are a way to gain control in advance and reduce uncertainty, framing the interaction with GenAI chatbots as a conversation or interview where rules should be clarified before making decisions. In contrast, they described just-in-time as delayed and disruptive (Section 5.2, 6.2, 7.3.1, and 7.4.1), while the timing of conversations is especially critical [36, 41]. This aligns with earlier findings by Balebako et al. [5], which showed that although participants remembered notices shown after app use, they considered them poorly timed for making meaningful privacy decisions. They also viewed that just-in-time disrupts the natural flow of the conversation between users and GenAI ecosystems, which is typically expected to be smooth and uninterrupted [52]. Yang et al. emphasize that both pragmatic and hedonic qualities shape users’ affective experience with conversational agents. Pragmatic qualities like fluidity and seamlessness support a natural and efficient interaction, while hedonic aspects such as comfort and technological pride enhance emotional engagement [52]. As a result, just-in-time privacy interfaces can inadvertently undermine the overall user experience in conversational contexts.

8.2 Unpacking Participants’ (Conflicting) Perceptions of First-party and Third-party Ecosystems

Our results suggested that participants did not hold a consistent privacy attitude toward a particular type of ecosystem. They perceived a lack of transparency in GenAI systems regarding how the chatbot ecosystem operates. As a result, participants formed a range of assumptions about first-party GenAI based on prior experiences, such as assuming the system is already connected, generates profit, or reuses information. Their perceptions also differed across the two scenarios, with these contextual differences shaping contrasting views in the Personalization Scenario and Data-sharing Scenario (Section 7.3.2 and 7.4.2).

In the Personalization Scenario, participants generally viewed the interaction as remaining within the same GenAI ecosystem. Gemini amplified perceptions of loss of control by raising concerns about ecosystem-wide data harvesting, as participants believed the

first-party company (i.e., Google) could access extensive user information. Participants noted that distributing data across soiled services could reduce exposure to any single entity (Section 7.3.2 and 7.4.2), consistent with prior work [42]. Similarly, prior work shows that when sharing data with extended service providers in GenAI ecosystems, participants tended to trust third-party ecosystems more, preferring not to put all their information concentrated within a single first-party ecosystem [47].

In the Data-sharing Scenario, participants emphasized the change of environment when data moved outside the original system (Section 7.5). They held strong assumptions about how first-party GenAI ecosystems operate, particularly regarding data usage, transfer, and profiling, while having limited knowledge about third-party GenAI systems (Section 7.3.2 and Section 7.4.2) and expressed greater needs for detailed privacy information when interacting with them (Section 6.3). Participants felt a greater sense of privacy control within familiar first-party ecosystems, reflecting their long-standing trust in these providers, as noted in prior usable privacy and security research. In this scenario, the first-party ecosystem enhanced perceived control by framing data sharing as a curated, accountable process within a trusted boundary, unlike external companies that seemed less accountable. Participants recognized that the institutional scale of companies like Google could provide stronger safeguards for their data.

Additionally, even though our study only focuses on two key factors and studies how they impact users' privacy perceptions, our participants, when thinking through various contextual factors, also considered the costs and benefits (Section 7.4.3). This notion resembles the literature in behavioral economics, suggesting that privacy decision-making is driven by a rational calculus of costs and benefits, also by misperceptions, emotional responses, social norms, and cognitive heuristics [1]. Future work may continue to explore how other behavioral contextual factors (e.g., company brand, etc.) and economic factors (e.g., the value of users' privacy, costs and benefits, etc.) may influence users' privacy perceptions, expectations, and behaviors in GenAI ecosystems.

8.3 Design Implications

Based on our findings, we offer several design recommendations to enhance users' sense of privacy control, reduce their concerns, and align users' privacy experiences with the unique characteristics of GenAI ecosystems.

8.3.1 Reconsider the Timing of Privacy Interfaces for GenAI Ecosystems. Designing usable privacy for the GenAI ecosystem necessitates a fundamental rethinking of timing strategies established for traditional websites and mobile applications. All studies show that just-in-time privacy choice (appearing at the time of data sharing) erodes trust, diminishes perceived control, and heightens transparency concerns. GenAI systems are designed for fluid, natural, human-like dialogue. Even though interruption privacy mechanisms are fairly common in most websites or mobile apps (e.g., notifications about sensor access in iOS), a similar approach in the GenAI ecosystem may risk users' experiences and degrade some key elements for adoption, such as conversational fluidity, user comfort, and confidence. Thus, having privacy choice at setup appears to be a better solution based on our study, which is also the

approach that ChatGPT is currently taking, yet their interfaces only include a very brief notice without offering users the ability to make changes. Additionally, future research may explore new paradigms in designing privacy choice for GenAI ecosystems. Designers could leverage native conversational mechanics, such as contextual turn-taking, agent-initiated clarification, or minimal confirmations, to surface privacy choice, rather than simply transplant other existing design alternatives (e.g., banners) [49].

8.3.2 Increase transparency to Enable Informed Privacy Decisions. GenAI ecosystems should increase transparency to help users build a clear and mature mental model, enabling informed privacy decisions. Our findings show that participants often filled gaps in knowledge with assumptions, both positive and negative, about how first-party and third-party ecosystems manage their data. In the Personalization Scenario, when inputting information into the GenAI chatbot, participants wanted clarity about the scope of data access, who would receive the data, and how the system would use it for personalization. Lacking such details, they often worried about ecosystem-wide data harvesting, particularly when the first-party company already possessed extensive user information. In the Data-sharing Scenario, when data is moved to external service providers, participants placed even greater weight on knowing which providers could access their data and how sharing occurred outside the original system. Because of vague mental models, they expressed stronger needs for transparency and more concrete privacy information. Clarifying assumptions about first-party ecosystems and making third-party practices explicit can reduce uncertainty, counter harmful assumptions, and ultimately reinforce user trust in both types of GenAI ecosystems. These concerns parallel the Digital Markets Act (DMA), which requires gatekeepers to obtain explicit consent before "cross-use" of personal data, such as combining data across services or enforcing cross-platform sign-ins. Participants' worries about first-party ecosystems reflect the high risks the DMA seeks to mitigate, while their call for explicit information about external providers aligns with the law's focus on informed consent and accountability.

8.3.3 Balancing Anticipatory and Contextual Control in Privacy Timing. Our results show that at-setup options that give users anticipatory control, reduce uncertainty, and promote natural interaction styles. Yet, they also have drawbacks. Participants noted that at-setup choices often feel intrusive, detached from specific tasks, and create cognitive overload at the beginning of the interaction. Just-in-time options address these gaps by providing task-specific relevance and contextual clarity, enabling users to make informed decisions when it matters most. However, they risk interrupting the natural conversational flow. A balanced design should therefore combine both approaches. First, a clear at-setup privacy choice can help users understand the overall data practices and establish expectations. Second, privacy choice should clarify the scope of that consent for specific actions or data flows, ensuring users know exactly what is shared, when, and with whom. This layered strategy can strengthen user trust without undermining usability.

8.3.4 Support Minimal and Task-oriented Data Use. Our participants showed a strong task-oriented preference in data sharing.

This behavioral pattern implies that GenAI ecosystems may implement granular, scenario-driven data governance triggered by concrete tasks (e.g., accessing calendar availability for scheduling or location for navigation). This highlights the importance of minimal, purpose-driven data practices. GenAI ecosystems could adopt granular, scenario-driven governance where access is triggered only by concrete needs (e.g., calendar availability for scheduling, location for navigation). By aligning data access with immediate user goals, systems not only reduce privacy concerns but also strengthen perceptions of fairness and necessity. Clear boundaries around “why this data is needed now” can reinforce trust and complement broader at-setup consent with precise, just-in-time justification.

8.3.5 Provide Effortless Data Management. When the GenAI ecosystem requests data for personalization, participants preferred options that offered ease of use and control. They expressed a strong desire for mechanisms that support effortless data selection, and the ability to review, delete, or retrieve shared information (Section 5.2 and Section 6.2). To complement these real-time and effortless controls, as Section 5.2 and 6.2 discussed, a dashboard that tracks data flow and allows users to manage data may be an ideal design. Similar to activity logs in mobile platforms or smart devices [45], it could include functions like revoking access, flagging questionable events, or receiving summaries of usage trends.

8.4 Limitation and Future Work

First, our work examines participants’ perceived satisfaction and preferences regarding different privacy choice timings, and perceived convenience or usefulness does not necessarily equate to alignment with users’ true intentions or best interests. An important next step is to assess whether such reported satisfaction genuinely reflects users’ underlying intentions and long-term privacy interests. Additionally, in practice, the mechanisms could be misused to reduce friction and encourage disclosure rather than to support informed and autonomous decisions. Future work could examine ethical and policy guidelines to mitigate these risks. Second, our study examined only a limited scope of the privacy design space, focusing on the timing of privacy choices. Building on this foundation, future research may explore other design dimensions of privacy notice and choice, including modality, type, functionality, and channel, to develop a more holistic understanding of privacy-related interactions.

Our vignette and storyboard design necessarily simplified interactions into a single decision point, and selected flight booking as the research context. The storyboard approach does not fully capture the multi-turn, evolving nature of real-world GenAI conversations. In practice, just-in-time privacy notices may emerge at different stages depending on users’ goals, the system’s inferences, or changing data needs across turns. Privacy preferences are also highly context-dependent, and users may respond differently to privacy interfaces in other scenarios. Future research could move beyond vignette-based simulations to more realistic, interactive settings, capturing multi-turn, dynamic interactions that may reveal how privacy choice timing and other design spaces of privacy notice and choice work in practice.

The marginal trends observed in our results highlight potential patterns that warrant further investigation in studies, providing

opportunities to refine design implications and deepen understanding of user privacy behaviors in GenAI ecosystems. Finally, all of our participants came from the US. We did not consider cultural or other differences (e.g., legislation) that may contribute to participants’ privacy perceptions and behaviors. Future work could extend this research to more diverse contexts.

9 Conclusion

We studied users’ privacy perceptions in GenAI ecosystems through two survey studies and a follow-up interview. Across personalization (Personalization Scenario) and data sharing (Data-sharing Scenario), both ecosystem type and timing of privacy choices shaped perceived control, trust, and concerns. Participants favored at-setup options for anticipatory control and reduced uncertainty, but also valued just-in-time options for task-specific clarity. Privacy attitudes toward first- vs. third-party ecosystems varied by context. In the Personalization Scenario, first-party systems raised concerns about ecosystem-wide data harvesting. In the Data-sharing Scenario, their scale and familiarity fostered accountability compared to unfamiliar third-party providers. These findings reflect contextual integrity: users judge privacy not by ecosystem type alone but by data type, purpose, and flow. Our work informs the design of privacy choices in GenAI ecosystems.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback and all participants for their participation. This work is in part supported by the National Science Foundation CNS-2341187, CNS-2426397, CNS-2442221, and a Google PSS Faculty Award.

References

- [1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (Jan. 2015), 509–514. doi:10.1126/science.aaa1465
- [2] A. Acquisti and J. Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE Security and Privacy Magazine* 3, 1 (Jan. 2005), 26–33. doi:10.1109/MSP.2005.22
- [3] Mutahar Ali, Arjun Arunasalam, and Habiba Farrukh. 2025. Understanding Users’ Security and Privacy Concerns and Attitudes Towards Conversational AI Platforms. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 298–316. doi:10.1109/SP61157.2025.00241
- [4] Shirley Anderson and Yuanfei Zhao. 2025. Generative AI Interface Design Considerations for Private Equity. In *Companion Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25 Companion)*. Association for Computing Machinery, New York, NY, USA, 51–55. doi:10.1145/3708557.3716352
- [5] Rebecca Balebako, Florian Schaub, Idris Adjerd, Alessandro Acquisti, and Lorrie Cranor. 2015. The Impact of Timing on the Salience of Smartphone App Privacy Notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (Denver, Colorado, USA) (SPSM '15)*. Association for Computing Machinery, New York, NY, USA, 63–74. doi:10.1145/2808117.2808119
- [6] Glen Berman, Nitesh Goyal, and Michael Madaio. 2024. A Scoping Study of Evaluation Practices for Responsible AI Tools: Steps Towards Effectiveness Evaluations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 294, 24 pages. doi:10.1145/3613904.3642398
- [7] Travis Breaux (Ed.). 2020. *An introduction to privacy for technology professionals*. International Association of Privacy Professionals, Portsmouth, NH.
- [8] Rune Haubo Bojesen Christensen. [n. d.]. Cumulative Link Models for Ordinal Regression with the R Package ordinal. ([n. d.]).
- [9] Rune Haubo Bojesen Christensen. 2010. ordinal: Regression Models for Ordinal Data. (March 2010). doi:10.32614/CRAN.package.ordinal Institution: Comprehensive R Archive Network Pages: 2023.12-4.1.
- [10] Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Stanton. 2022. Is it a concern or a preference? An investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. In

- Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. USENIX Association, Boston, MA, 331–346. <https://www.usenix.org/conference/soups2022/presentation/colnago>
- [11] Luca Deck, Jakob Schoeffer, Maria De-Arteaga, and Niklas Kühl. 2024. A Critical Survey on Fairness Benefits of Explainable AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 1579–1595. doi:10.1145/3630106.3658990
 - [12] Zane Durante, Qiuyuan Huang, Naoki Wake, Ran Gong, Jae Sung Park, Bidipta Sarkar, Rohan Taori, Yusuke Noda, Demetri Terzopoulos, Yejin Choi, Katsushi Ikeuchi, Hoi Vo, Li Fei-Fei, and Jianfeng Gao. 2024. Agent AI: Surveying the Horizons of Multimodal Interaction. doi:10.48550/arXiv.2401.03568 arXiv:2401.03568 [cs].
 - [13] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 342, 11 pages. doi:10.1145/3613905.3650750
 - [14] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (CHI '21). Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3411764.3445148
 - [15] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 64, 16 pages. doi:10.1145/3411764.3445148
 - [16] Abenezzer Golda, Kidus Mekonen, Amit Pandey, Anushka Singh, Vikas Hassija, Vinay Chamola, and Biplab Sikdar. 2024. Privacy and security concerns in generative AI: a comprehensive survey. *IEEE Access* 12 (2024), 48126–48144.
 - [17] 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>
 - [18] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. “Okay, whatever”: An Evaluation of Cookie Consent Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (CHI '22). Association for Computing Machinery, New York, NY, USA, 1–27. doi:10.1145/3491102.3501985
 - [19] 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 (2022), 59:1–59:42. doi:10.1145/3512906
 - [20] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. 2021. Toggles, Dollar Signs, and Triangles: How to (In)Effectively Convey Privacy Choices with Icons and Link Texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 63, 25 pages. doi:10.1145/3411764.3445387
 - [21] Akhil Jain, Vasistha Ved, ANISH Khadare, Jainam Chheda, Sarthak Deshmukh, and DN JAIN. 2024. Gen AI: A Survey of Security and Privacy Threats. *NFSU Journal of Cyber Security Digital Forensic* 3, 1 (2024), 39–43.
 - [22] Haojin Jin, Boyuan Guo, Rituparna Roychoudhury, Yaxing Yao, Swarn Kumar, Yuvraj Agarwal, and Jason I. Hong. 2022. Exploring the Needs of Users for Supporting Privacy-Protective Behaviors in Smart Homes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 449, 19 pages. doi:10.1145/3491102.3517602
 - [23] Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. Deduplicating Training Data Mitigates Privacy Risks in Language Models. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, Baltimore, Maryland, USA, 10697–10707. <https://proceedings.mlr.press/v162/kandpal22a.html>
 - [24] Patrick Gage Kelley, Lucian Cescas, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing privacy notices: an online study of the nutrition label approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10). Association for Computing Machinery, New York, NY, USA, 1573–1582. doi:10.1145/1753326.1753561
 - [25] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). Association for Computing Machinery, New York, NY, USA, 3393–3402. doi:10.1145/2470654.2466466
 - [26] Siwon Kim, Sangdoon Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. ProPILE: probing privacy leakage in large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, 20750–20762.
 - [27] Jan H. Klemmer, Stefan Albert Horstmann, Nikhil Patnaik, Cordelia Ludden, Cordell Burton, Carson Powers, Fabio Massacci, Akond Rahman, Daniel Votipka, Heather Richter Lipford, Awais Rashid, Alena Naiakshina, and Sascha Fahl. 2024. Using AI Assistants in Software Development: A Qualitative Study on Security Practices and Concerns. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*. Association for Computing Machinery, New York, NY, USA, 2726–2740. doi:10.1145/3658644.3690283
 - [28] Hao-Ping (Hank) Lee, Lan Gao, Stephanie Yang, Jodi Forlizzi, and Sauvik Das. 2024. “I don’t know if we’re doing good. I don’t know if we’re doing bad”: investigating how practitioners scope, motivate, and conduct privacy work when developing AI products. In *Proceedings of the 33rd USENIX Conference on Security Symposium (SEC '24)*. USENIX Association, USA, 4873–4890.
 - [29] Hao-Ping (Hank) Lee, Yu-Ju Yang, Thomas Serban Von Davier, Jodi Forlizzi, and Sauvik Das. 2024. Deepfakes, Phrenology, Surveillance, and More! A Taxonomy of AI Privacy Risks. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (CHI '24). Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3613904.3642116
 - [30] Haoran Li, Dadi Guo, Wei Fan, Mingshi Xu, Jie Huang, Fanpu Meng, and Yangqiu Song. 2023. Multi-step Jailbreaking Privacy Attacks on ChatGPT. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 4138–4153. doi:10.18653/v1/2023.findings-emnlp.272
 - [31] Tianshi Li, Sauvik Das, Hao-Ping (Hank) Lee, Dakuo Wang, Bingsheng Yao, and Zhiping Zhang. 2024. Human-Centered Privacy Research in the Age of Large Language Models. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (CHI EA '24). Association for Computing Machinery, New York, NY, USA, 1–4. doi:10.1145/3613905.3643983
 - [32] Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. 2022. Large Language Models Can Be Strong Differentially Private Learners. doi:10.48550/arXiv.2110.05679 arXiv:2110.05679 [cs].
 - [33] Jialiu Lin, Shahriyar Amini, Jason I. Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and purpose: understanding users’ mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. Association for Computing Machinery, New York, NY, USA, 501–510. doi:10.1145/2370216.2370290
 - [34] Pierre Lison, Ildikó Pilán, David Sanchez, Montserrat Batet, and Lilja Øvrelid. 2021. Anonymisation Models for Text Data: State of the art, Challenges and Future Directions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 4188–4203. doi:10.18653/v1/2021.acl-long.323
 - [35] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhamidi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 27–41. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>
 - [36] Xingyu Bruce Liu, Shitao Fang, Weiyan Shi, Chien-Sheng Wu, Takeo Igarashi, and Xiang ‘Anthony’ Chen. 2025. Proactive Conversational Agents with Inner Thoughts. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 184, 19 pages. doi:10.1145/3706598.3713760
 - [37] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (Nov. 2019), 23 pages. doi:10.1145/3359174
 - [38] Daniel Reinhardt, Johannes Borchard, and Jörn Hurlienne. 2021. Visual Interactive Privacy Policy: The Better Choice?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–12. doi:10.1145/3411764.3445465
 - [39] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security (SOUPS '15)*. USENIX Association, USA, 1–17.
 - [40] Sakib Shahriar, Sonal Allana, Seyed Mehdi Hazratifard, and Rozita Dara. 2023. A survey of privacy risks and mitigation strategies in the artificial intelligence life cycle. *IEEE Access* 11 (2023), 61829–61854.
 - [41] Gabriel Skantze. 2021. Turn-taking in Conversational Systems and Human-Robot Interaction: A Review. *Computer Speech & Language* 67 (May 2021), 101178. doi:10.1016/j.csl.2020.101178
 - [42] Daniel J. Solove. 2013. Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126 (May 2013), 1800–1903. Issue 7. <https://harvardlawreview.org/print/vol-126/introduction-privacy-self-management-and-the-consent-dilemma/>
 - [43] Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. 2024. Beyond Memorization: Violating Privacy Via Inference with Large Language Models. doi:10.48550/arXiv.2310.07298 arXiv:2310.07298 [cs].

- [44] Jingjing Sun, Jingyi Yang, Guyue Zhou, Yucheng Jin, and Jiangtao Gong. 2024. Understanding Human-AI Collaboration in Music Therapy Through Co-Design with Therapists. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 704, 21 pages. doi:10.1145/3613904.3642764
- [45] Parth Kirankumar Thakkar, Shijing He, Shiyu Xu, Danny Yuxing Huang, and Yaxing Yao. 2022. "It would probably turn into a social faux-pas": Users' and Bystanders' Preferences of Privacy Awareness Mechanisms in Smart Homes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 404, 13 pages. doi:10.1145/3491102.3502137
- [46] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Info. Sys. Research* 22, 2 (June 2011), 254–268. doi:10.1287/isre.1090.0260
- [47] Xingyi Wang, Xiaozheng Wang, Sunyup Park, and Yaxing Yao. 2025. Mental Models of Generative AI Chatbot Ecosystems. In *Proceedings of the 30th International Conference on Intelligent User Interfaces* (Italy) (IUI '25). Association for Computing Machinery, New York, NY, USA, 1016–1031. doi:10.1145/3708359.3712125
- [48] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. Ethical and social risks of harm from Language Models. doi:10.48550/arXiv.2112.04359 arXiv:2112.04359 [cs].
- [49] Ziang Xiao, Tiffany Wenting Li, Karrie Karahalios, and Hari Sundaram. 2023. Inform the Uninformed: Improving Online Informed Consent Reading with an AI-Powered Chatbot. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 112, 17 pages. doi:10.1145/3544548.3581252
- [50] Ziang Xiao, Q. Vera Liao, Michelle Zhou, Tyrone Grandison, and Yunyao Li. 2023. Powering an AI Chatbot with Expert Sourcing to Support Credible Health Information Access. In *Proceedings of the 28th International Conference on Intelligent User Interfaces* (IUI '23). Association for Computing Machinery, New York, NY, USA, 2–18. doi:10.1145/3581641.3584031
- [51] Yuting Xie, Meiyang Li, Siye Zeng, Jiaqi Mo, Yuting Diao, and Guan hong Liu. 2024. FitPal: Reshape Daily Exercise Misconceptions Among Elders through AI Chatbot and Community-based Services. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing* (CSCW Companion '24). Association for Computing Machinery, New York, NY, USA, 387–392. doi:10.1145/3678884.3681880
- [52] Xi Yang, Marco Aurisicchio, and Weston Baxter. 2019. Understanding Affective Experiences with Conversational Agents. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300772
- [53] Chao Zhang, Xuechen Liu, Katherine Ziska, Soobin Jeon, Chi-Lin Yu, and Ying Xu. 2024. Mathemyths: Leveraging Large Language Models to Teach Mathematical Language through Child-AI Co-Creative Storytelling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 274, 23 pages. doi:10.1145/3613904.3642647
- [54] Shikun Zhang, Lily Klucinec, Kyerra Norton, Norman Sadeh, and Lorrie Faith Cranor. 2024. Exploring Expandable-Grid Designs to Make iOS App Privacy Labels More Usable. In *Twentieth Symposium on Usable Privacy and Security* (SOUPS 2024). USENIX Association, Philadelphia, PA, 139–157. <https://www.usenix.org/conference/soups2024/presentation/zhang>
- [55] Zhiping Zhang, Michelle Jia, Hao-Ping (Hank) Lee, Bingsheng Yao, Sauvik Das, Ada Lerner, Dakuo Wang, and Tianshi Li. 2024. "It's a Fair Game", or Is It? Examining How Users Navigate Disclosure Risks and Benefits When Using LLM-Based Conversational Agents. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 156, 26 pages. doi:10.1145/3613904.3642385

A Appendix

A.1 Survey 1

Your Prolific ID:

Have you used any of the following AI products? (e.g., ChatGPT, Google Gemini) [Select all that apply]

- ChatGPT
- Google Gemini
- Others
- No (If selected, end the survey)

In which scenarios do you typically use AI chatbots? [Select all that apply]

- Work-related tasks
- Education/learning
- Daily life
- Entertainment/Social
- Health advice
- Sensitive tasks
- Other

Thank you for filling out the first part of the survey. We will now present four storyboards about using AI chatbots. Please answer the questions based on each storyboard.

Main Survey for Study 1

— Storyboard Here —

In the following interface, ChatGPT would like to use your information for personalized services.

Which option(s) would you add to ChatGPT for personalization? [Select all that apply]

- Memory in ChatGPT
- Google Account
- Microsoft Account

Why would you choose this/these option(s)?

What types of data do you think are useful for personalization? [Select all that apply]

- Calendar
- Email
- Contacts
- Photos
- Others

Why are these types of data useful for personalization?

What types of data do you think are too sensitive to use? [Select all that apply]

- Calendar
- Email
- Contacts
- Photos
- Others

Why are these types of data too sensitive?

Does this interface give you privacy control over your personal information?

Low, Somewhat Low, Neutral, Somewhat High, High

How has this interface changed your trust in ChatGPT's ability to handle your data responsibly?

- I trust it significantly more
- I trust it somewhat more
- No change
- I trust it somewhat less
- I trust it significantly less

How has this interface changed your willingness to continue using ChatGPT?

- I'm significantly more willing to use it
- I'm somewhat more willing to use it
- No change
- I'm somewhat less willing to use it
- I'm significantly less willing to use it

What kinds of controls would you like to have when ChatGPT requests access to your data?

- Choose which types of data ChatGPT can access
- Set time limits for data access
- Review/delete ingested data
- Real-time approval before transfer
- Automatic expiration of data
- Dashboard to track data sharing
- Other

What information is most important to understand when ChatGPT requests your data?

- Specific types of data being accessed
- Purpose of use
- Duration of storage
- Whether data is shared
- Who can access the data
- Security measures
- Legal compliance
- Data storage location
- Risks of exposure
- Contact for concerns
- Other

What are your main concerns when ChatGPT requests access to your data?

- Unauthorized data collection
- Misuse by AI
- Lack of transparency
- Sharing with untrusted third parties
- Loss of control
- Inability to track usage
- Other

Main Survey for Study 2 — Storyboard Here — Please carefully review the storyboard below. It illustrates a continuous scene that you should immerse yourself in. Imagine yourself as part of this scenario and respond to the following questions based on what you experience in the scene.

Low, Somewhat Low, Neutral, Somewhat High, High

How has this interface changed your trust in Gemini's ability to handle your data responsibly?

- I trust it significantly more

- I trust it somewhat more
- No change
- I trust it somewhat less
- I trust it significantly less

How has the interface changed your willingness to continue using Gemini?

- I'm significantly more willing to use it
- I'm somewhat more willing to use it
- No change
- I'm somewhat less willing to use it
- I'm significantly less willing to use it

Ideally, what kinds of controls would you like to have when Gemini notifies you about sharing data with third-party services (e.g., Expedia)? [Select all that apply.]

- Choose in advance which types of data Gemini can or cannot access (e.g., block location sharing)
- Set time limits for data access (e.g., "only use my data for 24 hours")
- Review/delete ingested data later
- Real-time approval before each data transfer
- Automatic expiration of shared data
- A dashboard to track who received my data
- Other (Please specify) _____

In this scenario, which of the following information is most important for you to understand when Gemini notifies you to share data with third-party services (e.g., Expedia)? [Select all that apply.]

- Specific types of data being accessed (e.g., emails, calendar entries, files in Drive)
- Purpose of using the data (e.g., personalization, model training, service improvement)
- How long will the data be stored (e.g., temporary vs. permanent storage)
- Whether data will be shared with third parties (e.g., advertisers, partner companies)
- Who can access the data (e.g., only Gemini, Google, Google Ecosystem)
- Security measures to protect the data (e.g., encryption, access controls)
- Legal compliance with privacy regulations (e.g., GDPR, CCPA)
- Physical location of data storage (e.g., servers in specific countries)
- Potential risks of data exposure (e.g., breaches, misuse)
- Direct contact for privacy concerns (e.g., email, support team)
- Other (Please specify) _____

When Gemini notifies you that it will connect to third-party services (e.g., Expedia), what are your main concerns? [Select all that apply.]

- Unauthorized data collection
- Potential misuse by the AI
- Lack of transparency about what's collected
- Data shared with untrusted third parties
- Loss of control over shared data
- Inability to track how data is used

- Other (Please specify) _____

Demographic Survey

1. How do you describe yourself?

- Female
- Male
- Non-binary
- Prefer not to answer

2. How old are you?

- 18–24
- 25–34
- 35–44
- 45–54
- 55+
- Prefer not to answer

3. What is your highest education level? [Select one]

- High school or below
- Bachelor degree
- Master degree
- Doctoral degree
- Prefer not to answer

4. What best describes your employment status over the last three months?

- Working full-time
- Working part-time
- Unemployed
- Stay-at-home parent
- Student
- Retired

5. Which best describes your comfort level with computing technology?

- Ultra Nerd: Build computers, run servers, code apps
- Technically Savvy: Go-to person for tech help
- Average User: Know enough to get by
- Luddite: Use only when necessary

A.2 Codebook for Survey

Table 7: The Codebook for Survey

Categories	Themes	Codes
Motivation to personalization (N=381)	Task-oriented temporal support (N=360)	Historical data for current task (N=150)
		Anticipatory planning (N=90)
		Time arrangement (N=120)
	Process the task (N=21)	Make the task easier (N=21)
Rules for choosing data sources (N=2459)	Process the task (N=211)	Scenario-based (N=211)
	Using patterns (N=671)	Personal preference (N=58)
		More data (N=613)
	Company (N=116)	Company/platform trust (N=90)
		Company/platform distrust (N=26)
	Convenience (N=78)	Easier (N=21)
		Already have data in memory (N=32)
		Same company (N=11)
		Easier to retrieve (N=14)
	Privacy control (N=1383)	Minimize data exposure (N=1208)
		Privacy Manage (N=64)
		Same company (N=2)
		Avoid others' info (N=85)
		Don't care (N=24)
Concerns choosing data source (N=712)	Potential risk (N=132)	Traditional privacy risks (N=111)
		AI-specific concerns (N=21)
	Private information (N=580)	Data is private (N=580)