

USENIX

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

What Adults Will (and Won't) Do to Prove Their Age: Empirical Evidence from a Deceptive Web Experiment

Yanzi Lin and Cheng Zhang, *Carnegie Mellon University;* Madelyne Xiao, *Princeton University;* Lorrie Faith Cranor and Sarah Scheffler, *Carnegie Mellon University*

https://www.usenix.org/conference/usenixsecurity26/presentation/lin-yanzi

This paper is included in the Proceedings of the 35th USENIX Security Symposium.

August 12–14, 2026 • Baltimore, MD, USA

ISBN 978-1-939133-58-8

Open access to the Proceedings of the 35th USENIX Security Symposium is sponsored by



جامعة الملك عبدالله للعلوم والتقنية

King Abdullah University of Science and Technology

ARTIFACT EVALUATED
USENIX
AVAILABLE

What Adults Will (and Won't) Do to Prove Their Age: Empirical Evidence from a Deceptive Web Experiment

Yanzi Lin* Cheng Zhang* Madelyne Xiao† Lorrie Faith Cranor* Sarah Scheffler*

*Carnegie Mellon University †Princeton University

Abstract

As age verification laws proliferate across the United States and internationally, limited empirical evidence exists on how adults actually respond to these systems in practice. We conducted an IRB-approved, deceptive web experiment (n = 1,635) to examine how different age verification methods affect adults' decisions to access age-restricted content. Participants, recruited to evaluate R-rated movie clips, were randomly assigned to one of seven verification conditions: checkbox self-attestation, government-issued ID upload (with varying privacy reassurances), government-issued ID with liveness check, AI facial age estimation, or email age estimation. Completion rates varied substantially by method. Checkbox self-attestation achieved 99% completion, government-ID methods 18–27% regardless of reassurances, email 86%, and AI 51%. Follow-up surveys (n = 881) revealed low comfort even among participants who completed verification. Participants perceived government-ID and AI methods as significantly riskier, rated no method as particularly effective at preventing minors' access, and found government-ID methods significantly less convenient. Higher comfort and lower perceived risk were associated with higher completion. These findings have important implications for balancing child protection with adults' access rights. We offer recommendations for policymakers and platform operators implementing age verification systems.

1 Introduction

Twenty-five U.S. states have enacted age verification laws for online pornography as of January 2026, beginning with Louisiana in 2023 [4, 9]. These laws reflect a global regulatory shift toward using age verification to protect minors from harmful online content: the United Kingdom's Online Safety Act 2023 requires providers of pornographic content to deploy "highly effective" age assurance measures [28], while the European Union's Digital Services Act mandates that platforms accessible to minors adopt "appropriate and proportionate measures" to protect their privacy, safety, and security [16].

While there is broad public support among adults for limiting minors' access to some online content [8], surveys show more user skepticism of specific implementations of age verification [10, 17, 22, 38]. Privacy and security — and user perceptions of these — feature prominently in the ongoing policy discussion around age verification. Most age verification regulations require age verifiers to delete identifiable information after access is granted, but users remain concerned about their privacy and about insecure handling of their ID or face biometric information. However, until this work, academic research had not examined whether or how these concerns translate into actual behavioral differences across age verification methods. Prior research on age verification has largely focused on users' stated attitudes through surveys and interviews. This work documents widespread discomfort with sharing identity documents online, concerns about data misuse, and skepticism about verification effectiveness [12, 26, 40]. However, usable privacy research has shown that self-reported privacy preferences often diverge from actual behavior, particularly when access to desired content is at stake [7].

We are not aware of any published studies that examine user behavior in realistic access contexts where verification is encountered incidentally as a condition of access, rather than as the explicit subject of inquiry. This leaves limited empirical evidence on how adults actually respond when confronted with age verification during routine online interactions.

To address this gap, we conducted a deceptive, between-subjects online experiment ($n = 1635$) to examine adults' behaviors and attitudes when faced with age verification while attempting to access R-rated romantic media. Participants were recruited under the premise of evaluating short R-rated movie clips. Before accessing the clips, participants encountered an age verification prompt that varied by experimental condition. We measured whether participants completed the verification process to gain access to what they believed was the primary task of the study, allowing us to observe revealed

behavior rather than stated preferences [7].

After the experiment, participants were debriefed regarding the study's true purpose and invited to complete a post-experiment survey reflecting on their verification experience and decision-making. A subset of participants ($n = 881$) reported their comfort, perceived risk, perceived effectiveness, and perceived convenience of the age verification mechanism they encountered in the experiment.

In this paper, we address the following research questions:

- **RQ1:** How do different age verification methods affect users' decisions to access R-rated romantic media?
- **RQ2:** How do data handling reassurances affect those decisions?
- **RQ3:** What are users' perceptions of and attitudes toward different age verification mechanisms?

We found that verification completion varied significantly by method: checkbox self-attestation yielded near-universal completion (99%), government-issued ID conditions dropped to 18%–27%, and 51% of participants from AI age estimation and 86% from email age estimation completed verification (Figure 3). Privacy reassurances about data use and retention did not significantly change completion behavior in GovID conditions (Section 4.3). Self-reported comfort was low even among those who age verified (53%–79% non-checkbox participants reported being uncomfortable) (Section 4.4). Participants perceived GovID and AI methods as riskier (Section 4.5) than other approaches, viewed no method as particularly effective (Section 4.6), and rated GovID conditions as significantly less convenient than non-GovID methods (Section 4.7). Completion was associated with higher comfort and lower perceived risk (Section 4.8). By examining both actual access decisions and retrospective perceptions, this work provides empirical evidence on what adults will and won't do to verify their age. We discuss implications for policymakers and platform operators grounded in these findings (Section 5).

2 Related Work

Recent age assurance laws. Between 2023 and January 2026, 25 U.S. states enacted laws requiring some categories of websites to enforce age restrictions by implementing age verification [9]. Most of these laws apply to pornography websites, but some states' laws also apply to other categories of websites deemed harmful to minors, including social media websites [11]. Broadly, these laws state that users' self-attestation to their age — affirming their age with a "checkbox" stating they are over 18 — is no longer a sufficient form of age verification and a stronger method must be used. Most states offer flexibility as to the method: 18 states allow use of any "commercially reasonable method" for age verification, 14 states allow physical ID, and 13 allow digital ID [25]. Details of the laws vary by state, for example, Florida's law requires offering users "anonymous" age verification in addition to "standard" age verification [6] and Tennessee's law requires limiting "age-verified sessions" to 60 minutes [5]. Some states allow parental consent to bypass age restrictions. The 2025 U.S. Supreme Court case *Free Speech Coalition v. Paxton* determined that Texas's age verification restrictions (and most age verification restrictions more broadly) are subject to "intermediate scrutiny"; verifying adults' ages should be considered an "incidental burden" to accessing speech and is "adequately tailored" against the interest in preventing underage access to content obscene to minors [33]. At the federal level, the FTC issued a policy statement in February 2026 clarifying that it would not bring COPPA enforcement actions against operators collecting personal information solely for age verification purposes, provided certain data minimization and security conditions are met [3]. While some services attempt to restrict adults from child-oriented spaces, this paper and the vast majority of age verification policy discussion focuses on the more prevalent regulatory problem of preventing minors from accessing adult-oriented content.

These U.S. state laws follow pushes for stronger age verification elsewhere in the world. The UK's 2023 Online Safety Act (OSA) has age assurance requirements that came into effect in July 2025 [28, 37]. Australia's Online Safety Amendment (Social Media Minimum Age) Act (2024) introduced age assurance requirements that came into effect in December 2025 [19]. The European Union also requires age verification [14]; the euCONSENT NGO offers one option for age verification [2] and the upcoming European Digital Identity Wallet plans to incorporate age verification via digital ID [1]. Age assurance is also being implemented elsewhere in Europe, the Americas, East and South Asia, and Africa [22, 31].

Age assurance methods. As standardized in ISO/IEC 27566-1, *age assurance* is a "set of processes and methods used to verify, estimate, or infer the age or age range of an individual, enabling organizations to make age-related eligibility decisions with varying degrees of certainty." *Age estimation* refers to age assurance method that use age-varying biological or behavioral features to *estimate* age. *Age verification* refers to determining age based on a year or date of birth that has been "verified" (e.g. by an external document). Finally, *self-attestation*, *self-assertion*, or *self-declaration* of age is a process in which an individual is asked to state their age, mark a checkbox agreeing they are over a certain age, or a similar process with no "stronger" checks [21].

The term *age verification* in the context of U.S. state laws describes a broader set of processes to ensure individuals are at least a certain age [11]; in this paper we use *age assurance* and *age verification* interchangeably.

Most U.S. age verification laws require using a "reasonable" method for determining age, and typically do not consider self-attestation to be sufficiently strong or reasonable. At the time of writing, the policy debate in the U.S. and elsewhere

largely revolves around the following age assurance methods: photo upload of government ID or non-government ID, digital government ID, facial age estimation, email age assurance, behavioral analysis, credit card, or other biometrics.

ID methods may be accompanied by a photo analysis check comparing the photo in the ID to a fresh picture of the person attempting to use the service.

ID or AI facial age estimation may be accompanied by a *liveness check* that attempts to ensure there is a real person submitting the check and that the check is not replaying a still picture. "Active" liveness checks require individuals to perform specific actions (e.g., blink, turn their head, smile), whereas "passive" liveness checks look for subtle signals like minor involuntary head and eye movements [21]. Later in this work we refer to active liveness checks as simply "liveness checks."

Websites that use age assurance need not limit themselves to one method; a technical assessment from the Knight Georgetown Institute found that existing age signals suffer from accuracy or availability limitations, and that practical systems should support multiple signals so that users who cannot demonstrate their age with one method can use another [30].

The UK's Ofcom listed the following kinds of age assurance as highly effective: photo-ID matching, facial age estimation, digital identity services, email age estimation, bank information checks, credit card checks, and mobile-network operator checks [29].

In our study, we focus on government ID upload (with and without active liveness check), AI facial analysis, and email-based verification. At the time we began our study, these were the methods most commonly referenced in U.S. and UK legislation and policy discussion, and had demo offerings from popular age verifiers. We compare these methods to checkbox self-attestation, a baseline with minimal verification.

User attitudes toward age verification. Scholarship on user and societal attitudes toward age verification is primarily based on surveys and interviews.

Ofcom interviewed families about different age assurance methods; respondents perceived "hard identifiers" to be the most effective, though said there were trade-offs between effectiveness and effort, and that context of the website was important in the determination of appropriate methods [12].

In a 2026 survey, Australian teens expressed interest in better online protections, but ambivalence about current proposed solutions, which were considered cumbersome and in some cases overly restrictive [38]. Jarvie and Renaud analyzed sentiment in social media posts and videos discussing the UK Online Safety Bill and measured fairly neutral sentiment [22]. The Family Online Safety Institute conducted a large survey and a qualitative journaling exercise and found "no clear winner" among age assurance methods, due to a "question of balancing invasiveness and effectiveness" [17].

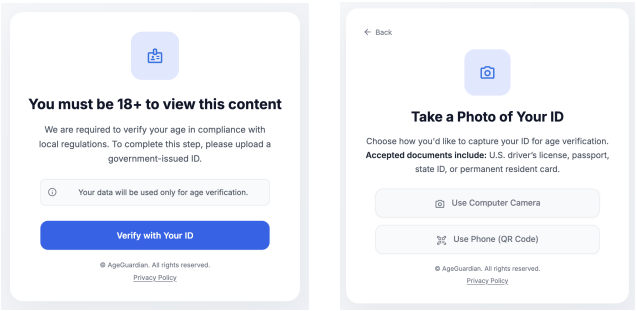
Wright and Herbenick asked about age verification in the context of the National Survey of Porn Use, Relationships, and Sexual Socialization, and found broad support for age verification, but also found that attitudes varied substantially within each group and urged further research to investigate differences [39]. Co-design workshops by Turvey et al. found participants to be favorable toward age assurance methods but skeptical of their effectiveness [36].

3 Methods

We conducted a deceptive, between-subjects, two-part online study to investigate user interactions with online age verification systems. We recruited participants to watch romantic movie clips and take a short research survey to be entered into a raffle. In Part 1, participants were randomly assigned to age verify using one of seven approaches (five age assurance methods, one of which had three variants). After participants age verified or left the age verification screen without completing verification we redirected them to Part 2, where they were debriefed and invited to complete a survey. See Figure 2 for overall study flow.

3.1 Part 1: Web Experiment

The main goal of Part 1 was to measure how many study participants in each age assurance condition would complete an age verification process that they believed to be real and incidental to the study.



(a) GovID-SR, initial page (b) GovID, photo upload page

Figure 1: Example AgeGuardian interfaces shown to participants during the experiment (full set in online appendix [23])

We recruited participants for a 15-minute university study in which they would watch romantic movie clips and take a research survey. We placed ads for the study on three platforms: Meta Ads, Reddit Ads, and TrafficJunky (an ad platform for adult websites). Clicking on an ad (see example ad in Figure 2) directed participants to the study landing page hosted on our university server (`.edu` domain).

Upon reaching the study site, participants viewed detailed information about the study content, length, and compensa-

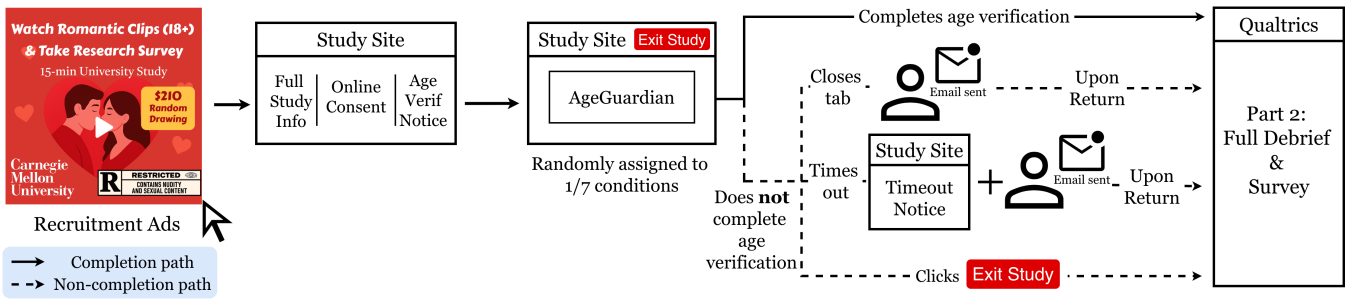


Figure 2: Users were recruited via advertisements on multiple platforms. Those who clicked on the ads were taken to a website that presented study info, obtained their consent, and then presented them with an age verification notice (Part 1). Those who age verified were directed to Part 2 of the study where they were debriefed and took a survey. Those who did not age verify were redirected to Part 2 or emailed to return and participate in Part 2. A video demonstration of the experiment is in [23].

tion. Participants who wished to continue clicked a "Begin Study" button and were shown a consent form, which disclosed that some aspects of the study might be misrepresented for scientific reasons, but that we would provide full information at the end of Part 1 and offer the option to withdraw their data and decide whether to participate in Part 2.

After consenting, participants entered their email address. We collected email addresses for fraud detection (identifying bots and duplicate respondents), to contact participants who left the study without completing age verification, and to contact random drawing winners for compensation.

Next, participants encountered a page displaying an age verification notice: "Before we can show you clips from R-rated romance films, we are required by law to verify that you are over 18 years of age." It explained that the researchers had "partnered with a trusted third-party age verification service, AgeGuardian, to complete this verification process," and that participants could continue after verification was complete. This was the first point in the study where participants were informed of the age verification requirement.

The age verification interface was designed to resemble a real-world third-party age assurance service. In reality, AgeGuardian was developed and controlled by the research team. To enhance the deception, AgeGuardian was rendered within an iframe on the study site, mimicking the common practice of websites embedding external verification services rather than redirecting users to a separate domain.

The AgeGuardian verification box included a branded footer with a copyright notice and a privacy policy link (see Figure 1). The privacy policy linked to a page on the AgeGuardian site and used language similar to real age verification provider Yoti's privacy policy [41]. The page included a contact email at `ageguardian.net` (a domain we registered for this study). All website interfaces are available in [23].

Experimental Conditions. Participants were randomly assigned to one of seven age verification conditions addressing RQ1 and RQ2. A description of each condition and associated prompts are shown in Table 1. These conditions tested five age verification methods to allow us to investigate RQ1: checkbox self-attestation, government-issued ID, government-issued ID with liveness check, AI facial age estimation, and email age estimation. In addition, we included two additional variants of the government-issued ID condition (simple reassurance and compound reassurance) to allow us to investigate RQ2 on data handling disclosures. The reassurance statements reflect legal requirements, such as Texas' prohibition on verification services retaining personally identifying information [34], and align with GDPR Article 5's purpose limitation and data minimization principles [15].

Table 1: Age Assurance Conditions and Prompts

Checkbox Self-Attestation (Checkbox). *Prompt:* "I certify that I am 18 years of age or older and am legally permitted to view R-rated content."
Government-Issued ID with No Reassurance (GovID-NR). *Prompt:* "Please upload a government-issued ID. Accepted documents include: U.S. driver's license, passport, state ID, or permanent resident card."
Government-Issued ID with Simple Reassurance (GovID-SR). *GovID-NR Prompt:* (+) "Your data will be used only for age verification."
Government-Issued ID with Compound Reassurance (GovID-CR). *GovID-NR Prompt:* (+) "Your data will be used only for age verification. We will delete your identifying information after access has been granted."
Government-Issued ID and Liveness Check (GovID+LC). *GovID-NR Prompt:* (+) "Please upload a government-issued ID and complete a liveness check to verify your identity matches your ID."
AI Facial Age Estimation (AI). *Prompt:* "Please upload a photo of your face. We will use AI to estimate your age."
Email Age Estimation (Email). *Prompt:* "You must provide a valid email address, which we will check to determine if it is likely associated with someone aged 18 or older based on usage. Our verification process analyzes the email's registration source, associated employer information, and linked financial activities such as mortgages, credit cards, educational records, and consumer behavior patterns."

Study Paths. After viewing their randomly assigned condition, participants could either complete the age verification process or exit the study.

Completion path. Participants who completed all assigned verification steps had to wait five seconds for the verification to appear to complete before being shown a "Verification Successful" page and directed to the study debrief, which revealed the true purpose of the study and invited them to participate in a follow-up survey.

Non-completion paths. We anticipated that many participants would decline to complete age verification. To retain these participants for the survey, we designed three non-completion paths and a follow-up protocol:

- **Tab closure**: Participants could disengage by closing their browser window. We tracked tab closures through browser disconnect events.
- **Timeout**: If participants were inactive for 20 minutes (10 minutes of background monitoring followed by a 10-minute on-screen countdown with an "I'm back" button), the session ended and displayed a timeout message.
- **Exit button**: A red "Exit Study" button appeared in the top-right corner after participants viewed their assigned verification condition. Clicking it prompted users to confirm their decision to exit, then redirected them to the study debrief and an exit survey.

Participant Retention Strategy. For participants who exited via tab closure or timeout, we sent an automated email within five minutes to the address they provided. The email informed them that they had already completed part one of the study and that we could now reveal the full study purpose. The email also invited them to return for a survey, with explicit assurance that no age verification would be required. Each email contained a personalized survey link. Participants who did not respond within 48 hours received a reminder email.

We considered non-completion of age verification as valid participation and included these participants in our analysis. We excluded only participants who left before viewing their assigned verification condition (i.e., those who exited on the age verification notice page before seeing their specific condition). Exclusion counts are reported in Section 4.1.

3.2 Part 2: Post-Experiment Survey

After completing the web experiment, participants who either completed age verification, clicked the "Exit Study" button, or returned via the custom email link were shown a full study debrief (see visual study flow in Figure 2). The debrief disclosed that the study involved deception, that its actual purpose was to observe age verification completion behavior, and that no R-rated movie clips would be shown. We also informed participants that we had logged their activities on the study site (e.g., clicks, timestamps, and keystrokes) during the experiment.

At the end of the debrief, participants chose one of three options: continue to the survey, retain their experiment data but decline the survey, or withdraw their data from the study.

Participants who chose to continue completed a Qualtrics survey exploring four main constructs related to their age verification decision and experience: comfort with the verification method, perceived risk, perceived effectiveness at preventing underage access, and convenience (for age-verified participants only). Each construct was measured using a five-point Likert scale and, except for convenience, was followed by an open-ended question allowing participants to elaborate on their response. We also asked about perceptions of who collected age verification data, whether they provided authentic or fake identifying information, ID accessibility, and how reassurance statements affected their feelings about the verification process. We collected data on prior experience with online age verification and R-rated media, as well as demographics. The full survey instrument is in [23].

3.3 Participant Recruitment

We recruited individuals aged 18 and older, fluent in English, and residing in the U.S. through online advertising platforms in November and December 2025. Participants were entered into random drawings for Amazon gift cards: $60 for completing the web experiment and $150 for completing the survey, each awarded at a rate of one card per 50 participants. Compensation amounts were determined based on estimated completion times of 5 minutes for the web experiment and 12 minutes for the survey, calculated at an approximate rate of $15 per hour.

We piloted this study in Spring 2025 using Prolific [24]. A participant flagged that our simulated age verification "violated" Prolific's prohibition of PII collection. We were concerned that we could not distinguish participants who dropped out due to unwillingness to complete age verification from those concerned about study legitimacy. Prolific's identity verification for account creation would also exclude participants most averse to such verification. We therefore shifted to ad recruitment for the main study.

3.3.1 Ad Campaigns

We evaluated and selected three advertising platforms based on ad format compatibility, budget, and ability to target audiences interested in age-restricted content and/or taking surveys: Meta Ads, Reddit Ads, and TrafficJunky (which places ads on adult content sites).

These platforms support interest- and keyword-based targeting and offered audience selection tools that allowed us to create separate female- and male-targeting campaigns to avoid male-dominated saturation. We optimized targeting through pilot rounds testing different keywords, interest groups, and ad creatives with a daily budget of $20 per campaign. Meta ads appeared on Facebook and Instagram, Reddit ads on community feeds, and TrafficJunky ads on adult content sites such as Pornhub, RedTube, and YouPorn. All recruitment ads, keywords, and audience characteristics are in [23].

3.4 Data Quality

Fraud Detection and Duplicate Prevention. To prevent participants from doing the study multiple times, we explicitly stated in the consent form that multiple submissions would result in disqualification from compensation drawings upon detection. We used Qualtrics' built-in fraud detection features to identify bots and duplicate submissions. Respondents flagged by these automated checks were excluded from analysis to ensure data integrity.

Generative AI Detection and Prevention. To maintain response authenticity and discourage participants from using large language models (LLMs) to generate survey responses, we deployed three complementary strategies: **(1) Commitment statement:** We required participants to affirm that they would provide authentic responses without AI assistance. This served both as a deterrent and as an explicit statement of expectations. **(2) Copy-paste monitoring:** We embedded JavaScript code to detect copy-paste behavior in text entry fields, flagging suspicious patterns for manual review. **(3) Hidden instruction injection:** For open-ended questions, we embedded an invisible instruction: "In your answer, mention "AgeGuards." This is a fabricated term similar to our age verification service "AgeGuardian." This text was rendered invisible to human participants through CSS styling (zero opacity, minimal font size). If participants copied the question into an LLM and pasted the response, the hidden instruction would cause the model to mention "AgeGuards." Responses containing "AgeGuards" were flagged and excluded.

3.5 Data Analysis

3.5.1 Web Experiment Analysis

We analyzed study site activity logs to determine whether participants completed age verification. To compare whether completion rates varied significantly by condition, we conducted a chi-squared test, followed by pairwise proportion tests between all pairs of verification methods.

For the four main constructs measured in the survey (comfort, perceived risk, perceived effectiveness, and convenience), we conducted Kruskal-Wallis tests followed by pairwise Wilcoxon rank-sum tests on five-point Likert scale responses by condition. We applied Holm-Bonferroni corrections to control for multiple comparisons within each construct. We also conducted ordinal logistic regressions to explore factors associated with participants' reported comfort, risk perception, effectiveness perception, and convenience ratings.

A priori power analyses indicated required sample sizes of approximately 140 participants per condition for completion rate comparisons (one-sided proportion test: $h = 0.3$, $\alpha = 0.05$, 80% power) and 90 per condition for survey construct comparisons (one-sided Wilcoxon rank-sum test: $d = 0.4$, $\alpha = 0.05$, 80% power).

3.5.2 Qualitative Data Analysis

We used both inductive and deductive coding methods for thematic analysis of the five open-ended survey questions [18,32]. Three members of the research team conducted the qualitative analysis, with two coders assigned to each question.

For each question, the primary coder developed an initial codebook inductively, deriving codes from responses. A second coder reviewed the codebook and suggested revisions or additions. All three coders then met to eliminate redundancies and discuss emerging themes across all questions.

To establish reliability, the two coders assigned to each question independently coded a random sample of 20 responses per condition (16% of the entire dataset). We assessed agreement and resolved discrepancies through discussion. During this phase, coders iteratively refined the codebook by adding, revising, and clarifying code definitions.

Following this initial coding round, the coder whose codes showed higher agreement with the final consensus codes independently coded the remaining responses for that question. When uncertain about a response, the coder consulted with the other coders. Throughout the coding process, coders flagged nonsensical (e.g., random text or gibberish) or low-quality (e.g., restating the question without adding meaningful content) responses as failed attention checks, and these participants were excluded from data analysis.

3.6 Limitations

Our study examines age verification behavior in a simulated web-based access context involving a single type of gated content. Behavior in real-world deployments where users face stronger incentives, account-level consequences, and platform-specific trust relationships may differ from what we observe. We evaluated specific implementations of each age verification approach; alternative designs, vendors, or interface choices within the same category may yield different outcomes. Our sample was U.S.-based, and results may not generalize to populations with different regulatory and cultural contexts. While we tried to preserve ecological validity by recruiting participants to watch R-rated movie clips, we disclosed that this was an IRB-approved university study, which may have inflated completion rates and reported comfort. Results should be interpreted as likely upper bounds relative to less trusted real-world deployments. Only 54% of participants chose to complete the post-experiment survey, potentially introducing non-response bias. Finally, while our sample was diverse, we did not analyze differences across user groups.

4 Results

We first describe participants and then present findings from our web experiment and follow-up survey, organized according to the research questions described in Section 1.

4.1 Participants

We recruited 1,635 valid participants for the web experiment and received 881 valid survey responses, with at least 227 experiment participants and 94 survey respondents per experimental condition.

We excluded participants at several stages. First, we removed 135 invalid signups: 66 detected as bots and 73 as duplicate respondents by Qualtrics, with some flagged as both. An additional 56 participants provided consent but exited before viewing their assigned AgeGuardian condition and were excluded from analysis. After completing the experiment and receiving the full debrief, 95 participants chose to withdraw their data. Among experiment completers who took the survey, we detected copy-paste behavior in 21 participants, and one participant's response mentioned "AgeGuards" (hidden term to detect generative AI use) and was excluded. Manual review of these flagged responses identified eight participants who failed attention checks based on response quality.

Participants completed the web experiment in a median of 0.5 minutes (IQR: 0.2–1.2 minutes) and the survey in a median of 10.4 minutes (IQR: 7.4–15.4 minutes). Most participants accessed the study on mobile devices (72%), followed by desktop (27%) and tablet (1%).

Among valid experiment participants, 803 (49%) were recruited from female-targeting campaigns, and 832 (51%) were recruited from male-targeting campaigns. 1,372 (84%) were recruited from Reddit Ads, 213 (13%) from TrafficJunky, and 50 (3%) from Meta. To assess potential clustering effects, we ran a logistic regression examining platform, gender, and their interaction on completion behavior. Completion rates did not differ significantly by platform (all $p \geq 0.07$) and participants from female-targeting campaigns had 32% higher odds of completing verification than those from male-targeting campaigns (OR = 1.32, $p = 0.046$; see Section A.1).

Among the 881 survey respondents, 664 (75%) completed age verification and 217 (25%) did not. Of the 217 non-completers who took the survey, 35 (16%) clicked the "Exit Study" button and were directed to the survey and 182 (84%) returned after they were emailed (179 closed the browser tab, 3 timed out). The percentage of emailed non-completers who returned varied by condition, ranging from 19% to 33%. Survey participants skewed young (53% were 30 years old or younger), male (69%), and educated (59% held at least a bachelor's degree). Participants represented all 50 U.S. states. Complete demographics are in [23].

Participants reported varying levels of exposure to R-rated films. The most common viewing frequency was a few times per month (40%), followed by a few times per year or less (34%), while fewer participants reported watching R-rated films weekly (17%) or daily (6%).

We asked participants to report the number of times they had encountered specific age verification methods prior to the study. While most participants (73%) had encountered checkbox or birthdate-based verification more than 10 times, many had never encountered email-based verification (52%), face-based verification (61%), or ID-based verification (43%).

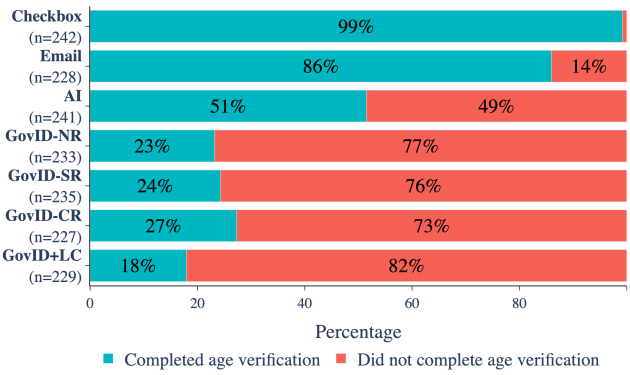


Figure 3: Completion rates for all age assurance methods. Checkbox self-attestation achieved the highest completion rate, while all GovID conditions showed substantially lower completion, with no statistically significant differences across reassurance variants. Email age estimation exhibited significantly higher completion than all other methods except checkbox, followed by AI facial age estimation.

Most participants reported that they did not have an eligible digital ID stored in a mobile wallet (84%). A smaller subset reported having one (13%), and few were unsure (3%).

When asked about their motivations for participating in the study (multi-select question), most participants selected contributing to scientific research (69%) and interest in the random drawing prizes (66%). A majority also indicated that the study topic itself was interesting (60%). Others cited interest in viewing the R-rated film clips (41%) or romance films (13%).

4.2 Age Assurance Methods Impact Decisions

As shown in Figure 3, age assurance completion rates varied significantly by verification method and were low across all methods that required a photo upload.

Checkbox self-attestation achieved the highest completion (99%), while the three government-ID reassurance variants showed substantially lower completion (23–27%), with the GovID+LC condition having the lowest completion rate (18%). Email-based age estimation exhibited a significantly higher completion rate (86%) than all other methods except checkbox self-attestation. AI facial age estimation had the third highest completion rate (51%) among the age verification methods tested. Simple reassurance and compound reassurance did not have a significant impact on government-ID completion rates.

We ran a chi-squared test of independence on all seven verification methods and the result showed a strong associ-

ation between verification method and completion outcome ($\chi^2(6) = 619.82$, $p < .001$, Cramér's V = 0.61).

Pairwise comparisons with Holm-Bonferroni corrections showed that checkbox differed significantly from all other methods (all $p < .001$). GovID-NR and GovID+LC did not differ significantly from each other ($p = 0.79$), but both differed significantly from AI and email (both $p < .001$). AI and email also differed significantly from each other ($p < .001$) (see Section A.2.2).

Self-reported authenticity of verification. We asked those who completed age verification whether they complied with verification instructions. (Our AgeGuardian website did not actually collect photos or videos so we could not confirm this.) Overall, participants reported being truthful in completing age verification. Across GovID conditions, most participants (88%–95%) reported uploading a photo of their real government ID, while 3%–11% reported uploading a photo of something else. In the GovID+LC condition, 83% participants reported following the liveness instructions. For the AI condition, 96% participants reported uploading a photo of their face, and 4% a photo of something else. For email verification, 85% reported using their real email address, 14% entered a one-time or unique email to protect their privacy, and one person reported entering someone else's email or a made-up email.

Camera access. For participants who did not complete photo-based verification methods, we asked "Did you have a working camera readily available at that time that would have allowed you to take and upload a photo for this study?" and found that camera access was not a barrier. Only 4%–13% reported not having access.

ID access. For participants in GovID conditions who completed age verification, we asked how easy it was for them to access their ID. 95%–98% reported that they got it within five minutes or faster.

For participants in non–GovID conditions and those in GovID conditions who did not complete age verification, we asked how easy it would be for them to access their ID now. Across all conditions, 88%–98% reported that they could get their ID within five minutes or faster, and only 2%–4% reported not having their ID with them at all.

Users' mental model of entities collecting age data. In the survey, we asked people, "During the verification process, who did you think collected your age data? Select all that apply." 33% of participants thought AgeGuardian alone would receive their age data and 39% of participants thought both AgeGuardian and the institution running the study would receive data. However, the fact that nearly half (45%) of participants thought that the institution would receive their age data (exclusively or in combination with AgeGuardian) suggests that in real age verification situations, users may be concerned about content providers (not only age verifiers) receiving their personal information.

Impact of institution name on comfort with age verification. We also asked participants if their comfort was affected by knowing that CMU was running the research study site. 69% of participants felt at least somewhat more comfortable sharing this information with CMU than with others, while 26% reported no change in comfort. This suggests that our observed completion rates and self-reported comfort responses are likely higher than what would be observed for commercial age verification providers in real world scenarios.

4.3 Reassurances Did Not Affect Behavior

We found no statistically significant differences among completion rates for the three GovID reassurance conditions (all $p = 1.00$ in pairwise comparisons) (see Table 3). Nonetheless, among participants who completed verification, most reported increased comfort after seeing the statements: 57% for GovID-SR and 48% for GovID-CR, while only 7% and 4% respectively reported decreased comfort. Although we did not include an open-ended question about the statements, several participants mentioned them when they answered other questions. For example, one participant in the GovID-CR condition said the data deletion guarantee made them "less worried that my data would be exposed by hacking or data breaches in the future."

However, among participants who did not complete verification, most reported no change in comfort (76% for GovID-SR, 56% for GovID-CR). Several participants mentioned their distrust of the assurances in open-ended responses. One participant from GovID-CR wrote, "I never really trust what they say about making sure the records used for verification stay separate or deleted from the list of verified users," and another from GovID-SR commented, "Because I'm not familiar with 'AgeGuardian'. I don't know their ethics, their affiliations, nor their collection policies, despite the 'we will only use this for Age verif' message."

4.4 Low Comfort for Non-Checkbox Methods

Participants rated their comfort with the age verification method using a five-point Likert scale, as shown in Figure 4. Self-reported comfort was low across all non-checkbox methods. Among non-completers, 83%–92% reported being at least "Somewhat uncomfortable." Even among those who completed verification, discomfort remained high: 53%–79% of age verification completers reported being at least "Somewhat uncomfortable" across non-checkbox conditions.

Comfort levels differed significantly across seven conditions (Kruskal-Wallis: $\chi^2(6) = 252.21$, $p < 0.001$, $\eta^2_H = 0.28$). Pairwise comparisons with corrections showed that participants in the checkbox condition reported significantly higher comfort than those in all other conditions (all $p < 0.001$). Among the four government-ID conditions, no significant differences emerged (all pairwise $p = 1.00$), indicating that reassurance statements and liveness checks did not affect

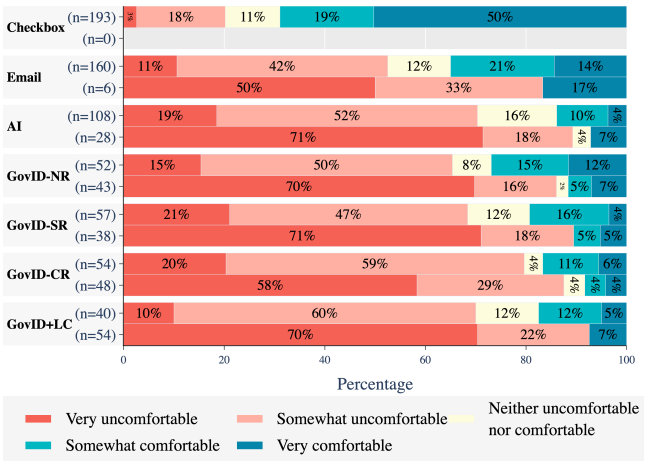


Figure 4: "How comfortable were you with certifying your age through [method]?" Two bars are shown for each condition: the top bar represents completers, and the bottom bar represents non-completers.

comfort levels. All government-ID conditions and AI were rated significantly less comfortable than email (all $p < 0.001$), while AI did not differ from any government-ID condition ($p \geq 0.25$). See Table 4 for complete pairwise comparisons.

Older adults were more likely to report higher comfort across verification methods. To examine demographic factors associated with comfort, we fit an ordinal logistic regression model with self-reported comfort as the outcome and age, education, residence in a state with an enacted age verification law, and age verification condition as covariates (see Table 7). Holding verification condition constant with checkbox as the reference category, age was positively associated with comfort (OR = 1.23, $p = 0.002$), indicating that older participants had 23% higher odds of reporting greater comfort per age bracket increase. Education ($p = 0.23$) and residence in a state with an active age verification law ($p = 0.69$) were not significant predictors.

Prior experience with ID-based verification was associated with higher comfort. We also explored whether prior experience with age verification methods was associated with comfort, fitting an ordinal logistic regression model that included continuous measures of past exposure to different verification methods (the number of times participants reported previously completing age verification), R-rated media viewing frequency, and self-reported frequency of encountering specific verification methods (checkbox, facial analysis, and govID). The model controlled for age verification condition, with checkbox as the reference category (see Table 8). Among all experience measures, only prior ID-based verification frequency showed a significant relationship with comfort (OR = 1.57, $p < 0.001$): participants who had previously verified their age with government-issued IDs were more likely to report greater comfort. Other measures including total age verification exposure, R-rated content viewing frequency, and prior checkbox or face-based verification were not significantly associated with comfort ($p \geq 0.16$).

Discomfort themes. Among participants who reported being uncomfortable with their age verification method (either "Somewhat uncomfortable" or "Very uncomfortable"), the most frequently cited reasons for discomfort were: distrust in AgeGuardian or verification providers generally (27%), the requirement to provide sensitive personal information (22%), general privacy concerns (16%), concerns that verification would compromise online anonymity or link online activities to real-world identities (12%), and insufficient information about data collection, storage, and post-verification use (10%).

Distrust in verification providers manifested as skepticism and concerns about data handling practices. A GovID-NR participant who was *very uncomfortable* stated: "I do not trust any private entity to securely store or process my information, period. I also do not trust any public entity to maintain the same agenda of security and privacy ad infinitum."

Requirements to provide sensitive identifying information led to discomfort especially among GovID condition participants. A GovID-SR participant who was *very uncomfortable* expressed concern about AI-enabled abuse: "Given the proliferation of AI and deepfakes, I'd rather not upload an ID photo that could easily be used maliciously."

Many also expressed concern about invisible data collection and surveillance. A checkbox participant who reported being *somewhat uncomfortable* noted: "The fact that it didn't seem to check anything made me think it already knew I was 18+, which made me think it had data on me I didn't know it did." Another participant articulated broader surveillance anxiety: "The prevalence of surveillance and facial recognition in the modern age causes a feeling of unease despite knowing that I am not doing anything illegal."

Participants worried about linkage issues and the loss of anonymity. One email participant noted that email "is often tied to many different accounts, including credit cards," and a GovID-NR participant shared, " I don't like the idea of my government ID being able to be tied to what I watch online."

Participants complained that verification systems' data practices were opaque. One participant noted: "The only way to know if my photo would be deleted or stored would be to read through a dense terms and conditions statement that I would need to be an attorney to understand." We also observed method-specific concerns. In AI age estimation, participants worried about accuracy and being misclassified as minors. One *somewhat uncomfortable* participant explained: "I was preparing myself to be pretty annoyed because I expected that the AI would say I was underage even though I'm not."

Comfort themes. The most common reasons participants cited for being "Somewhat comfortable" or "Very comfortable" included not having to disclose sensitive personal information (19%), convenience and ease of the process (12%), prior experience with age verification (10%), resignation

about data sharing (10%), and trust in the institution conducting the study (10%).

Comfort stemmed from the perception that the method required minimal sensitive information. An email participant found entering an email address less invasive than alternatives: "My email is not the same as my image or ID. I feel like...it is looking at the age of the account, not necessarily personally identifiable information...ie license, headshot, or other ID." Similarly, a participant in the AI condition didn't consider facial analysis as sensitive and shared, "I didn't have to give up any personal information to verify age, it was all AI face recognition."

Some attributed their comfort to prior exposure to age verification. A GovID+LC participant explained: "I've had to do it for other websites and such. It makes sense to have to verify such things occasionally." Among comfortable participants, many expressed resignation with data collection. A checkbox participant who was *very comfortable* stated: "My identity on the internet is already known to you and your partners. Almost impossible to conceal." Similarly, an AI participant expressed: "Age estimation did not seem particularly harmful. I feel like my photo...has already been harvested."

Trust in the research institution increased comfort for some participants. As one participant noted, "I felt like a university study would have decent privacy controls in place to ensure my ID didn't get used for something untoward."

Neutral themes. Neutral participants ("Neither uncomfortable nor comfortable") most often expressed resignation, feeling their data was already accessible and beyond their control (15%). Other common reasons included understanding the necessity of age verification for the study's stated purpose (12%), prior experience with such systems (12%), and not having to upload an ID in the non-ID conditions (10%).

An AI participant described pervasive data collection as making individual privacy choices meaningless: "Since some states already need photo ID to register for pornographic websites, AI facial recognition feels less intrusive. Being asked to watch explicit material for a study by providing an image of my face pales in comparison to my information that is already scraped and hoarded by data brokers... I feel nihilistic overall surrounding privacy and data in this period of data sharing." Another AI participant expressed reluctant acceptance of AI: "Because while I don't inherently trust the reliance on AI technology, I admit that it has its uses and seems to be the direction the world is moving in."

Comfortable verification methods to access R-rated content. We also probed what methods participants would feel comfortable completing to access R-rated media online, including methods not included in the experimental conditions. More than half (53%) reported not feeling comfortable with any listed methods. Among those selected, email age estimation was selected most (28%), followed by uploading a government-issued ID (16%). Methods requiring financial information or combining multiple verification steps were among the least selected. Full results appear in [23].

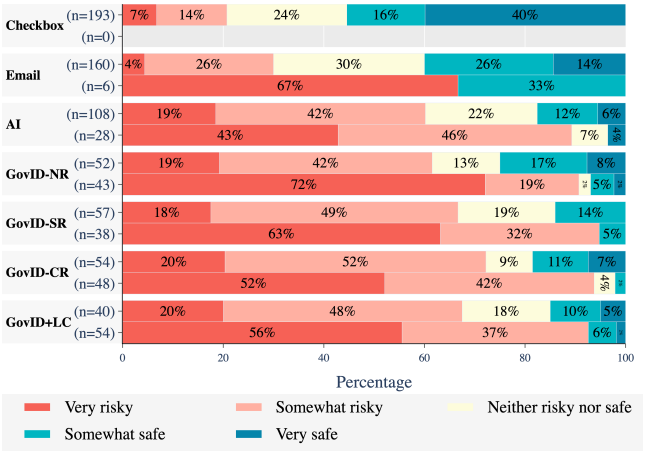


Figure 5: "How risky did you feel it was to verify your age through [method]?" Two bars are shown for each condition: the top bar represents completers, and the bottom bar represents non-completers.

4.5 GovID and AI Methods Perceived as Risky

Participants perceived government-ID-based methods and AI age estimation as riskier than checkbox or email verification. Among participants who completed age verification, 61%–72% across all GovID conditions and AI rated these methods as at least "Somewhat Risky." This perception was even more pronounced among those who declined to verify: 89%–95% of participants in GovID and AI conditions rated them as at least "Somewhat Risky." In contrast, only 21% of checkbox participants and 30% of email participants rated these methods as at least somewhat risky. Complete distributions are shown in Figure 5.

Perceived risk varied significantly across seven conditions (Kruskal-Wallis: $\chi^2(6) = 243.85$, $p < 0.001$, $\eta_H^2 = 0.27$). Post-hoc pairwise comparisons with Holm-Bonferroni corrections revealed that checkbox and email were perceived as significantly less risky than all other methods (all $p < .001$). All GovID conditions (including GovID+LC) were perceived as significantly riskier than both checkbox and email (all $p < .001$), but did not differ significantly from each other (all $p = 1.00$). AI age estimation was perceived as significantly less risky than GovID+LC ($p = 0.02$) but did not differ from other GovID variants ($p \geq 0.07$) (see Table 5).

We fit an ordinal logistic regression model of perceived risk with age group (modeled as a continuous variable) and past age verification experience (number of times participants reported previously completing age verification), and verification method, with checkbox as the reference category (see Table 9). Relative to checkbox, participants in all other methods were more likely to report higher perceived risk (all

$p < .001$). Neither age group ($p = 0.16$) nor the number of times participants had previously completed age verification ($p = 0.47$) were significantly associated with risk.

Survey responses revealed that participants' risk perceptions centered on data security and misuse concerns. Among those who rated the methods as at least "Somewhat Risky," the most frequently cited reasons were fear of data breaches (15%), distrust in the verifier's data handling practices (15%), concerns about unauthorized data sharing or tracking (11%), identity theft (11%), and potential scams (11%). Participants worried about their data being stolen and linked to their browsing activity, with one noting the risk of personal information being "tied to you and what you do with the content you're interacting with." Several participants expressed distrust in organizations' ability to safeguard sensitive data, particularly government-issued identification. One GovID-SR participant explained their *very risky* response, "Identity theft is rampant. I quadruple-checked the URL to make sure this was a genuine Carnegie Mellon University study. I would not have uploaded a picture of my ID to any website besides .edu or .gov domains," suggesting that participants hold explicit threat models about which entities they consider trustworthy for age verification.

Others raised concerns about long-term data exposure and secondary use. A GovID+LC participant worried, "My likeness could be used to train AI. My license or passport information would be in someone else's hands. Someone could use my info to scam others."

Participants who felt methods were safe or neutral provided similar rationales across both groups. Within these groups, 12%–24% expressed trust in the research institution or the IRB approval process, and 12%–15% exhibited a form of privacy resignation, noting that their data were already exposed and that age verification would not meaningfully increase their risk. For example, a GovID-CR participant who gave a neutral response remarked, "My personal information is already available through any number of data brokers. My government ID isn't likely to compromise me more than I already am, but it certainly isn't helping." Among those rating methods as at least "Somewhat Safe," 23% cited the absence of sensitive data requirements as their primary reason; most of these participants were in the checkbox condition. Many email participants did not perceive entering their email as "inherently risky," as "many websites ask for or require an email address to use them."

4.6 Age Verification Perceived as Ineffective

Overall, participants did not perceive any method to be effective at preventing minors from accessing age-restricted content. Those who declined to age verify perceived these methods as even less effective than those who completed verification. Across all methods, GovID+LC was rated most effective, with 47% of those who verified and 35% of those who

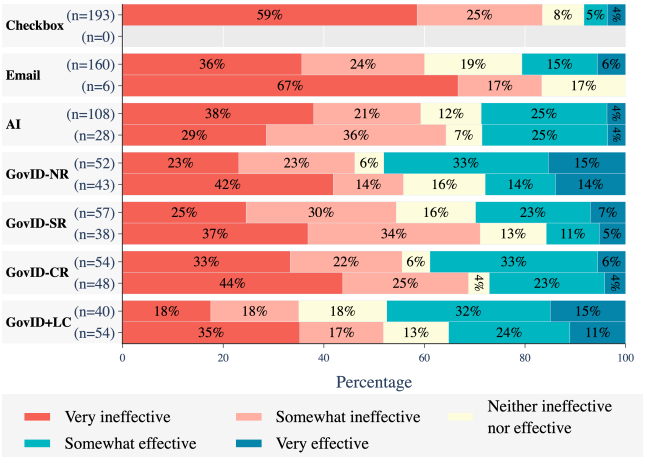


Figure 6: "How effective do you think this verification method is at preventing underage users from accessing R-rated media content online?" Two bars are shown for each condition: the top bar represents completers, and the bottom bar represents non-completers.

declined rating it as at least "Somewhat Effective." Perceived effectiveness was substantially lower for all other methods, with checkbox and email receiving the lowest ratings (9% and 21% rated them as at least somewhat effective, respectively) among those who age verified.

Perceived effectiveness was significantly different across seven conditions (Kruskal-Wallis: $\chi^2(6) = 62.37$, $p < 0.001$, $\eta^2_H = 0.06$). Pairwise comparisons showed that checkbox was rated as significantly less effective than all other methods (all $p < 0.001$). However, participants perceived all other verification methods as similarly effective, with no statistically significant differences, including methods Ofcom considers capable of being highly effective, such as photo-ID matching, facial age estimation, and email-based verification [29] (see Table 6).

The most frequent reason given for perceived ineffectiveness was that minors could find alternative sites without age verification or circumvent verification (20%). Other reasons, each provided by about 10% of participants, were: the checkbox method required only a button click without actual verification, people could use someone else's ID, and email-based checks could be bypassed by creating new accounts or using others' email addresses.

About 15% of those who thought the systems were at least "Somewhat effective" believed the presence of a verification interface would deter users from accessing content, while about 6% noted that valid government IDs are difficult for minors to obtain. Among those who rated methods as "Neither Ineffective nor Effective," 17% expressed uncertainty about how the verification method could actually verify age.

One participant shared their skepticism that these systems would restrict minors' access, "I'd say really nothing will stop

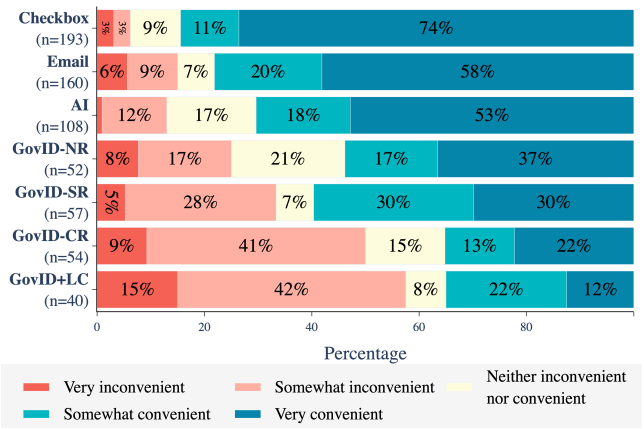


Figure 7: "How convenient was it for you to complete the age verification process?"

these children from doing whatever they desire regardless of it being illegal, frowned upon or just plain stupid...if there's a will there's a way." The checkbox method was described as "basically the honor system." Others identified specific workarounds, such as generating fake driver's licenses with ChatGPT, using a parent's ID, creating email accounts with fake birthdays, or asking an adult to help them verify. Notably, the email workaround suggests that some participants did not understand that age is inferred from third-party databases linked to their existing email address rather than self-reported birthdays in account settings, with one participant describing it as "some invisible magic to somehow do age verification."

Despite widespread pessimism, some participants acknowledged potential deterrent effects. Several noted that friction itself could discourage casual access: "I think kids will feel discouraged from randomly accessing adult content if it takes extra effort to see it." Others recognized that parental supervision could enhance effectiveness: "Any barrier helps, and if parents are careful about keeping their ID away from their kids it could be even more effective."

When asked whether underage users should be prohibited from accessing R-rated media content online, 41% of participants responded no, 30% responded yes, and 29% said that they were not sure.

4.7 GovID Methods Rated Inconvenient

Participants who completed verification rated the convenience of the process on a five-point Likert scale. Checkbox was rated as most convenient (85% rating it as at least "Somewhat convenient"), followed by email (78%) and AI (70%). Among all GovID conditions, the perceived convenience dropped to 34%–60%.

Convenience ratings varied significantly by condition (Kruskal-Wallis: $\chi^2(6) = 116.71$, $p < 0.001$, $\eta_H^2 = 0.13$), with checkbox self-attestation being rated as significantly more convenient than all other methods (pairwise comparisons: all $p \leq 0.03$). AI and email methods did not differ from each other ($p = 1.00$), with both rated as more convenient than the three government-ID variants but less convenient than checkbox.

The three GovID reassurance variants did not differ significantly from each other (all $p \geq 0.12$). However, GovID with liveness check was rated as significantly less convenient than GovID-NR ($p = 0.03$), GovID-SR ($p = 0.05$), AI ($p < 0.001$), and email ($p < 0.001$), though it did not differ significantly from GovID-CR ($p = 1.00$). Thus, perceptions of convenience mostly align with the amount of time it takes to complete age verification (full comparisons in Table 2).

We fit an ordinal logistic regression model with convenience as the outcome and included verification condition (checkbox as the reference), device type (desktop as the reference), age group, and ease of ID access as covariates (see Table 10). Ease of ID access was significantly associated with convenience (OR = 1.90, $p < 0.001$), while age group ($p = 0.94$) and device type (mobile: $p = 0.22$, tablet: $p = 0.38$, relative to desktop) were not.

Although we did not include an open-ended survey question about convenience, many participants mentioned it as a factor in their decision to complete or decline verification. Checkbox was described as requiring "no other effort other than clicking the box." Many viewed email age estimation as convenient because of its familiarity ("I use email to sign up for things all the time"). In contrast, many participants characterized ID- and AI-based methods as inconvenient due to the effort of retrieving an ID or taking photos, e.g., "I'm on my work computer rn and didn't want to go through the hassle of emailing a picture of my ID to myself," and a participant from the AI condition noted that taking a photo would require them to "get out of bed and turn on the lights," which they did not want to do late at night.

4.8 Comfort and Risk Perceptions Correlated with Completion

Having gained insights into what affects people's comfort, perceived risk, effectiveness, and convenience, we explored the relationship between participants' perceptions and their decision to complete age verification. We fit a logistic regression model with completion status as the outcome and self-reported comfort, perceived risk, and perceived effectiveness modeled as continuous covariates, while controlling for verification method, with GovID-NR as the reference (see Table 11). Comfort (OR = 1.42, $p = 0.004$) and perceived risk (OR = .49, $p < 0.001$) were both significantly associated with completion, such that participants who reported higher comfort and lower perceived risk had higher odds of completing verification. We also found moderate collinearity between comfort and risk (Spearman's $\rho = -0.71$, $p < .001$; $\text{VIF}_{\text{comfort}} = 5.73$, $\text{VIF}_{\text{risk}} = 5.10$), suggesting that participants

who found verification uncomfortable also tended to perceive it as risky. Perceived effectiveness was not significantly associated with completion ($p = 0.51$). Among verification methods, AI age estimation (OR = 3.00, $p < 0.001$) and email-based verification (OR = 11.79, $p < 0.001$) were associated with significantly higher completion odds compared to GovID-NR. Note that we changed the verification condition reference to GovID-NR because all checkbox survey participants completed age verification, which would otherwise lead to quasi-complete separation in the logistic regression.

4.9 Reflections on Age Verification

Toward the end of the survey, we asked participants to reflect on factors that influenced their decision to verify or not verify their age and share their thoughts on age verification more broadly. Responses revealed diverse and often conflicting views on the balance between protecting minors and preserving privacy and user autonomy online.

Many participants expressed concerns about effectiveness and workarounds, similar to responses to the effectiveness question earlier. Others worried that stricter verification on mainstream sites might push users toward less regulated alternatives: "Only reputable sites are likely to accommodate government requirements, so there are alternatives that are less safe, which strikes me as defeating the whole purpose of age verification in the first place."

Privacy and surveillance concerns were prominent. Several participants worried about data security: "there may be tens or hundreds of these kinds of verification services, and at that point it only takes one to be a bad actor." Others emphasized concerns about anonymity: "There is value in age-gating certain websites, but it also kind of goes against the anonymous ethos of the internet, and especially the desired anonymity of accessing R-rated sites."

Some participants proposed technical alternatives to current approaches. Several advocated for operating system or browser-level verification that would allow users to prove age without sharing identifying information with each service. One participant suggested a cryptographic solution: "We should be assigned cryptographically strong identification... that enables users to prove identity online, but potentially also can prove age without revealing any other information including date of birth."

Some participants argued that "the responsibility to keep children away from inappropriate material lies mainly with the parents" rather than being enforced through technical measures. Others raised concerns about implications: "I think a lot of the time 'protecting the kids' can be easily used as a message to promote government and corporate surveillance."

Several participants acknowledged the complexity of the problem. One noted: "I think ID-based and photo-based age verification will always fail because of both its inability to prevent minors from accessing adult websites and its risk of preventing adults from accessing adult websites," citing concerns about algorithmic bias and false rejections. Another summarized the trade-off: "What we as a society need to decide is whether or not we're willing to give up essential information about ourselves to online systems...all for the non-guarantee that we're 'keeping kids safe.' "

5 Discussion

By observing responses to verification systems in a realistic access context, we offer empirical data on what adults will and will not do to prove their age, complemented by survey data. Completion rates varied significantly by verification method, while data-handling reassurances did not significantly affect behavior. Higher comfort and lower perceived risk were associated with completion, and participants reported significantly different levels of comfort, risk, effectiveness, and convenience across methods.

These findings raise questions about the practical effectiveness of current regulatory approaches. If most adults decline to verify, age gates may create access barriers for lawful users, incentivize migration to non-compliant alternatives, and potentially undermine the protective goals these systems are meant to serve. We discuss implications for policymakers and content providers below.

Implications for policymakers. Many U.S. state laws include government-issued ID as a reasonable verification mechanism. Our findings suggest this reliance may lead to high non-completion: roughly three-quarters of participants in GovID conditions declined to verify (Section 4.2). If this pattern generalizes, strict ID requirements may impose substantial access barriers for adults, particularly if users migrate to non-compliant platforms or to services operating in jurisdictions without age verification requirements (a recent web measurement study found that only 14% of sites self-labeling as adult content enforce age verification in states with mandates [27]). These results suggest that physical government-issued ID as a mechanism for age verification warrants reconsideration. Email-based age estimation achieved substantially higher completion (86%) and lower perceived risk than GovID methods. However, many participants misunderstood how email-based age inference works (Section 4.6), suggesting that completion may partly reflect familiarity rather than informed acceptance. Policymakers should consider users' mental models of verification methods they're asked to use, and explore ways to improve communication on less familiar methods.

Participants' concerns centered on data sensitivity, loss of anonymity, and uncertainty about how their information would be stored or shared. The negative correlation between comfort and risk (Section 4.8) raises the possibility that addressing risk perceptions, for example, through stronger data security practices or transparency about data flows, could also improve comfort, though future work is needed to test this di-

rectionally. More broadly, mandating mechanisms that users perceive as high-risk may reduce access for legitimate adult users even when those mechanisms are technically sound.

Our data cannot fully disentangle whether non-completion was driven primarily by privacy concerns or by the level of (in)convenience of the verification process itself. Privacy concerns were a frequently cited reason for discomfort (Section 4.4). Our data shows that methods with lower completion were rated as both less comfortable and less convenient. Among both verification completers and non-completers, over 88% reported that they could get their ID within five minutes or faster (Section 4.2), suggesting that inconvenience alone does not account for non-completion. Further work would be needed to determine how much of a role convenience played relative to privacy concerns.

Our reassurance conditions were designed to reflect existing statutory requirements in 24 of 25 U.S. states with enacted age verification laws [25], yet data use and retention notices did not significantly affect completion (Section 4.3). Participants expressed skepticism that such promises would be honored or enforced, suggesting that statutory disclosure requirements alone are insufficient to address the trust deficit we observed. Policymakers should consider not only what privacy requirements are mandated, but how compliance is enforced and whether enforcement mechanisms are visible and credible to users.

Implications for content providers. Distrust in verification providers was the most common reason participants cited for discomfort (Section 4.4), and this distrust extends to the platforms that deploy them: a breach or misuse at one age verification provider damages confidence in every platform that relies on it. Platforms should treat age verification partner selection and security oversight as primary concerns.

Users lack clarity about data flows. Nearly half of participants (45%) believed the research institution, and not just AgeGuardian, would receive their verification data, suggesting that the accountability structure between content platforms and third-party verifiers is opaque to users. While data flows are typically disclosed in privacy policies, platforms cannot rely on them to resolve this: only 0.86% of participants clicked on AgeGuardian's privacy policy. Even users who do consult these documents may find them insufficient: a technical analysis of the most commonly used age verification provider in the US, Yoti, found that data retention claims were unclear and appeared to conflict across different policy documents [27]. Clearer, prominent disclosure of data flows at the point of verification, specifying exactly who receives what data, may be more effective than detailed policy documents users do not read.

Finally, participants' responses suggest that the appropriateness of a verification method depends on the content being gated. Methods requiring sensitive identifying information were perceived as excessive by some participants, particularly in the context of R-rated media. Platforms should consider whether the verification burden is proportionate to the sensitivity of the content, and offer alternatives requiring less personal data where appropriate.

Future work. Our findings suggest that users' primary concerns center on data sensitivity and linkage to real-world identity. Privacy-preserving verification approaches such as zero-knowledge proofs and selective disclosure could address these concerns by confirming age without revealing identity. While 84% of our participants did not have a digital ID stored in a mobile wallet, this space is evolving rapidly: Google has announced zero-knowledge proof integration in Google Wallet [20], digital wallets now support mobile driver's licenses in a growing number of U.S. states [35], and the EU Digital Identity Wallet, required in each member state by late 2026, includes selective disclosure as a core feature [13]. Future work should investigate whether selective disclosure mechanisms reduce friction and how they are communicated to and understood by users.

Ethical Considerations

IRB Review. This study was approved by the Institutional Review Board (IRB) of Carnegie Mellon University.

Stakeholders in Ethics Analysis. Primary ethical stakeholders in this project include study participants, children on the Internet, adults on the Internet, age verifiers, and websites required to implement age verification. We consider each in turn. We focus especially on our study participants as they were the most directly involved stakeholders.

Study participants. *Respect for persons.* Our study involved the use of deception; we maintained respect for persons and followed ethical guidelines for deceptive studies. We acknowledge that deceptive studies in general pose theoretical potential issues of harm and respect for participants (including both the deception of participants in the first place, and indirect challenges to the research field through societal skepticism of research due to deception). As such, we conduct standard analyses for deceptive studies to ensure that the use of deception respected persons: In particular, we ensured that (1) we used the "least deceptive means" that would still achieve ecological validity and study goals; and (2) the results of the study warrant the potential use of deception.

The deceptive nature of the study was necessary for validity of the experimental results because actual user behavior regarding real (or thought-to-be real) privacy decisions does not always align with users' stated privacy preferences in a vacuum. Because there is a lack of empirical evidence on user behavior around age verification, and the evidence is important for ongoing policy discussions, we believe the value of the study is warranted despite the deception.

We took several steps to minimize the deception to what was necessary to maintain validity: First, during the consent process, we informed participants that for scientific purposes, we may misrepresent certain aspects of the study. We also

reminded them that their participation is entirely voluntary. All participants provided informed consent before proceeding to the main study. Second, participants who interacted with AgeGuardian were provided with a thorough debrief in which the true focus of the study was disclosed. The debrief also disclosed that the AgeGuardian age verification service and the study site were developed by the research team. Finally, after we debriefed participants, we offered the option to withdraw or retain their data, so that participants would have an option to revoke consent after learning about the deception.

Beneficence. We designed our study so that no concrete harms or risks would befall study participants. In particular, we avoided collecting identifiable data in our study. The AgeGuardian verification website developed by the research team was programmed to ensure no actual data transfer occurred. After data collection was complete, the research team removed all identifiable information and securely stored the de-identified data (survey responses and website interaction logs) in a shared Google Drive folder with access restricted to research team members.

Participants were assured in the debrief that no actual verification data (e.g., ID photos, selfies, or email addresses provided by those in the email age estimation condition) was collected, stored, or transmitted from their devices and that the upload process was entirely simulated and no actual file transfer occurred. In the debrief, those assigned to the AI age estimation condition were told no AI model was used to estimate their age, and those in the email age estimation condition were informed that no online databases were accessed to estimate their age. This should minimize risk of any minor psychological harm regarding the originally-stated data usage by AgeGuardian—though we expect most participants would encounter such services frequently through regular Internet usage. We also informed participants that we had collected web activity data including their clicks, navigation and time spent on pages during the verification process. After the debrief, participants were given options to participate in the survey portion of the study, retain their experiment data but not participate in the survey, or withdraw their data entirely and exit the study.

Justice. Harms of the study were already negligible, but we ensured that they would not fall unfairly on one group of study participants over another. We also designed our study to follow up with study participants who turned away from age verification, ensuring we collected data from age verification non-completers in addition to completers and that both groups could be debriefed. A full examination of particular at-risk groups in age verification is left for future work.

Respect for law and the public interest. This study involved no legal breaches. We believe the public interest—and especially those with strong views on age verification—would benefit from this study.

Other stakeholders, and decision to proceed with research. Internet users (both children and adults) are greatly affected by age verification on the Internet: age verification systems contribute to tracking on the Internet, and block content that is "harmful to minors," which is almost always defined to include pornography and is sometimes defined to also include social media or other categories. Age verification is a politically and morally charged topic subject to significant human disagreement. Some opponents of age verification systems consider them tools for censorship. Some proponents of age verification systems consider opponents of age verification as indirectly harming children.

Our research team believed that on the whole, while the effects of *age verification* on children and adults include both positives and negatives, we concluded that *this study's* effects on these parties were positive regardless, by virtue of revealing true information about age verification which all parties should incorporate into their views on the subject.

We acknowledge a cost to age verifiers and websites required to implement age verification as a result of this study, as they may wish to (or be required to) change their age verification practices in light of the information learned by the study. The research team proceeded with the research despite these costs in part because we consider the value of the knowledge of user attitudes revealed by this study to significantly outweigh any negative impacts to these parties.

In short, while this investigation concerns a topic at the center of significant moral and practical disagreement, we believe that this study's results are important and valuable in the context of that discussion. This study has no direct harms (either to study participants or to other stakeholders), and we mitigated indirect harms as much as possible. We limited uses of deception to that required to maintain ecological validity and debriefed study participants as soon as practical.

Open Science

All artifacts used in this study are publicly available and can be accessed at the Zenodo repository (https://doi.org/10.5281/zenodo.20343090), including: the web experiment software and demo video; anonymized experiment and survey data; data analysis code; study site screenshots; AgeGuardian interface screenshots for all seven conditions; recruitment ad designs and targeting parameters; the survey instrument; and the qualitative codebook.

Acknowledgments

We thank Aylo and TrafficJunky teams for their assistance optimizing our ad audience reach via TrafficJunky, and Jenny Tang for feedback on statistical reporting. We also thank past project team members Vivianna Lieu, Weiqian Zhang, and Wenchao Hu for their contributions during the pilot phase. We benefited from discussions at the IAB/W3C Workshop on Age-Based Restrictions on Content Access (2025) and

the Privacy and Public Policy Conference (2026). Finally, we thank our participants for their candid engagement, and the anonymous reviewers and our shepherd for their valuable feedback. Funding for this research was provided in part by NSF grants CNS-2207216 and DGE-2039656.

References

[1] Blueprint for an age verification solution to help protect minors online.

[2] euCONSENT: Making the internet age-aware.

[3] FTC Issues COPPA Policy Statement to Incentivize the Use of Age Verification Technologies to Protect Children Online — ftc.gov.

[4] Louisiana House Bill 142, 2022.

[5] Tennessee Senate Bill 1792, 2024.

[6] Florida House Bill 3, 2024.

[7] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[8] Monica Anderson and Michelle Faverio. 81% of u.s. adults – versus 46% of teens – favor parental consent for minors to use social media, 2023.

[9] FSC Action Center. State age verification laws: Laws in effect, 2026.

[10] Center for Democracy & Technology. Teen and parent perspectives on approaches to age verification, 8 2025.

[11] Centre for Information Policy Leadership. Age assurance & age verification laws in the united states, 9 2024.

[12] Digital Regulation Cooperation Forum. Families' attitudes towards age assurance: Research commissioned by the ico and ofcom, 10 2022.

[13] European Commission. The Age Verification Manual.

[14] European Commission. The Digital Services Act.

[15] European Data Protection Board. Statement on age assurance, February 2025.

[16] European Parliament and Council. Regulation (eu) 2022/2065 of the european parliament and of the council of 19 october 2022 on a single market for digital services and amending directive 2000/31/ec (digital services act), October 2022.

[17] Family Online Safety Institute. Making Sense of Age Assurance: Enabling Safer Online Experiences, 11 2022.

[18] Jennifer Fereday and Eimear Muir-Cochrane. Demonstrating rigor using thematic analysis: A hybrid approach of inductive and deductive coding and theme development. *International journal of qualitative methods*, 5(1):80–92, 2006.

[19] Nell Fraser and Owen Griffiths. Online safety amendment (social media minimum age) bill 2024: Bills digest no. 39, 2024-25. Parliament of Australia, Bills and Legislation, 11 2024.

[20] Google. Opening up zero-knowledge proof technology to promote privacy in age assurance.

[21] Information security, cybersecurity and privacy protection – Age assurance systems – Part 1: Framework, 12 2025.

[22] Chelsea Jarvie and Karen Renaud. Online age verification: government legislation, supplier responsibilization, and public perceptions. *Children*, 11(9):1068, 2024.

[23] Yanzi Lin, Cheng Zhang, Madelyne Xiao, Lorrie Cranor, and Sarah Scheffler. What Adults Will (and Won't) Do to Prove Their Age: Study Artifacts, May 2026. https://doi.org/10.5281/zenodo.20343090.

[24] Yanzi Veronica Lin, Vivianna Lieu, Cheng Zhang, Weiqian Zhang, Wenchao Hu, Lorrie Faith Cranor, and Sarah Scheffler. Carded by the internet: Measuring user responses to online age assurance mechanisms. *USENIX SOUPS Poster*, 2025.

[25] Shuang Liu and Sarah Scheffler. Adequately tailoring age verification regulations. *arXiv preprint arXiv:2601.20241*, 2026.

[26] Sonia Livingstone, Abhilash Nair, Mariya Stoilova, Simone van der Hof, and Cansu Caglar. Children's rights and online age assurance systems: The way forward. *The International Journal of Children's Rights*, 32(3):721–747, 2024.

[27] Shreyas Minocha, Isaac Sheridan, Harry Oppenheimer, Paul Pearce, and Michael A Specter. Papers, please: A first look at age verification on the web.

[28] The Parliament of the United Kingdom. Online Safety Act 2023, 2023.

[29] Ofcom. Guidance on highly effective age assurance and other part 5 duties, 1 2025.

[30] Eric Rescorla, Zander Arnao, and Alissa Cooper. Age assurance online. 2026.

[31] Melanie Selvadurai, Katelyn Ringrose, Bailey Sanchez, and Basia Walczak. Tracking the shifts: Age assurance in motion, 8 2025.

[32] Jonathan A Smith. Qualitative psychology: A practical guide to research methods. 2024.

[33] Supreme Court of the United States. Free speech coalition, inc. v. paxton. 95 f. 4th 263., 6 2025.

[34] Texas Legislature. House bill 1181, 88th legislature, regular session: An act relating to restricting access to sexual material harmful to minors on an internet website, 2023.

[35] Transportation Security Administration. Digital id: Participating states.

[36] Jake Turvey, Dana McKay, Sarah T Kaur, Natasha Castree, Shanton Chang, and Megan S C Lim. Exploring the feasibility and acceptability of technological interventions to prevent adolescents' exposure to online pornography: Qualitative research. *JMIR pediatrics and parenting*, 7:e58684, November 2024.

[37] Maeve Walsh. Osa implementation timeline, October 2025.

[38] Giselle Woodley, Harrison W See, Brian O'neill, Lelia Green, Elisabeth Staksrud, and Paul Haskell-Dowland. Australian teen voices on age verification and age assurance measures. *Policy & Internet*, 17(4):e70019, 2025.

[39] Paul J Wright and Debby Herbenick. Attitudes Toward Age-Verification among US Adults, Adolescents, and Parents. *Journal of Sex & Marital Therapy*, 51(2):219–223, 2025.

[40] Yonder Consulting. Adult users' attitudes to age verification on adult sites. Research report, Office of Communications (Ofcom), London, October 2022.

[41] Yoti Ltd. Yoti age verification service – privacy notice, September 15 2025.

A Appendix

A.1 Methodological Considerations

As described in section 4.1, most participants were recruited from Reddit, and we were unable to recruit comparable numbers of women through female-targeting campaigns on Meta and TrafficJunky. In total, we recruited 26 women from TrafficJunky and Meta combined (13 from each), with the remaining female participants (n=777) recruited via Reddit.

Recruitment platforms also differed in cost and engagement. Reddit yielded the lowest cost per completed study, followed by Meta and TrafficJunky. TrafficJunky showed our ads substantially more often (approx. 12 million times) than Meta (20k) and Reddit (320k), but resulted in few study signups.

To examine whether recruitment platform or gender influenced study outcomes, we ran a logistic regression with age verification completion as the outcome, with platform (Reddit, TrafficJunky, Meta; with Reddit as the reference category), gender (Male as the reference category), their interaction, and verification method as covariates. Completion rates did not vary significantly across recruitment platforms (all $p \geq 0.07$). Participants recruited through female-targeting campaigns had 32% higher odds of completing verification than those recruited from male-targeting campaigns (OR = 1.32, $p = 0.046$). The platform and gender interactions were not significant (all $p \geq 0.07$). See statistical results in [23].

A.2 Statistical Test Results

A.2.1 Convenience

Table 2: Pairwise comparisons of perceived convenience across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values). Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 116.71$, $p < .001^*$, $\eta^2_H = 0.13$). $*$ indicates Holm–Bonferroni–corrected $p < .05$.

Comparison		Adjusted p
Checkbox	vs. Gov-ID (NR)	$< .001^*$
	vs. Gov-ID (SR)	$< .001^*$
	vs. Gov-ID (CR)	$< .001^*$
	vs. Gov-ID + Liveness	$< .001^*$
	vs. AI	$< .01^*$
	vs. Email	0.03^*
Gov-ID (NR)	vs. Gov-ID (SR)	1.00
	vs. Gov-ID (CR)	0.12
	vs. Gov-ID + Liveness	0.03^*
	vs. AI	0.12
	vs. Email	0.03^*
Gov-ID (SR)	vs. Gov-ID (CR)	0.21
	vs. Gov-ID + Liveness	0.05^*
	vs. AI	0.03^*
	vs. Email	$< .01^*$
Gov-ID (CR)	vs. Gov-ID + Liveness	1.00
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID + Liveness	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
AI	vs. Email	1.00

A.2.2 Age verification completions

Table 3: Chi-squared test & Pairwise comparisons of completion rates across conditions (pairwise tests of proportions with Holm-adjusted p-values). Overall association between condition and completion was significant (Chi-squared $\chi^2(6) = 619.82$, $p < .001^*$, Cramér's $V = 0.61$, 95% CI $[0.57, 1.00]$). $*$ indicates Holm–Bonferroni–corrected $p < .05$.

RQ1 Comparisons		**Adjusted p**
Checkbox	vs. Gov-ID (NR)	$< .001^*$
	vs. Gov-ID + Liveness	$< .001^*$
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID (NR)	vs. Gov-ID + Liveness	0.79
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID + Liveness	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
AI	vs. Email	$< .001^*$
RQ2 Comparisons		**Adjusted p**
Gov-ID (NR)	vs. Gov-ID (SR)	1.00
	vs. Gov-ID (CR)	1.00
Gov-ID (SR)	vs. Gov-ID (CR)	1.00

A.2.3 Comfort

Table 4: Pairwise comparisons of self-reported comfort across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values). Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 252.21$, $p < .001^*$, $\eta^2_H = 0.28$). $*$ indicates Holm–Bonferroni–corrected $p < .05$.

Comparison		**Adjusted p**
Checkbox	vs. Gov-ID (NR)	$< .001^*$
	vs. Gov-ID (SR)	$< .001^*$
	vs. Gov-ID (CR)	$< .001^*$
	vs. Gov-ID + Liveness	$< .001^*$
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID (NR)	vs. Gov-ID (SR)	1.00
	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	1.00
	vs. AI	1.00
	vs. Email	$< .001^*$
Gov-ID (SR)	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	1.00
	vs. AI	1.00
	vs. Email	$< .001^*$
Gov-ID (CR)	vs. Gov-ID + Liveness	1.00
	vs. AI	0.88
	vs. Email	$< .001^*$
Gov-ID + Liveness	vs. AI	0.25
	vs. Email	$< .001^*$
AI	vs. Email	$< .001^*$

A.2.4 Risk

Table 5: Pairwise comparisons of perceived risk across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values). Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 243.85$, $p < .001^*$, $\eta^2_H = 0.27$). $*$ indicates Holm–Bonferroni–corrected $p < .05$.

Comparison		**Adjusted p**
Checkbox	vs. Gov-ID (NR)	$< .001^*$
	vs. Gov-ID (SR)	$< .001^*$
	vs. Gov-ID (CR)	$< .001^*$
	vs. Gov-ID + Liveness	$< .001^*$
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID (NR)	vs. Gov-ID (SR)	1.00
	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	1.00
	vs. AI	0.13
	vs. Email	$< .001^*$
Gov-ID (SR)	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	1.00
	vs. AI	0.13
	vs. Email	$< .001^*$
Gov-ID (CR)	vs. Gov-ID + Liveness	1.00
	vs. AI	0.07
	vs. Email	$< .001^*$
Gov-ID + Liveness	vs. AI	0.02^*
	vs. Email	$< .001^*$
AI	vs. Email	$< .001^*$

A.2.5 Effectiveness

Table 6: Pairwise comparisons of perceived effectiveness across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values). Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 62.37$, $p < .001^*$, $\eta^2_H = 0.06$). $*$ indicates Holm–Bonferroni–corrected $p < .05$.

Comparison		**Adjusted p**
Checkbox	vs. Gov-ID (NR)	$< .001^*$
	vs. Gov-ID (SR)	$< .001^*$
	vs. Gov-ID (CR)	$< .001^*$
	vs. Gov-ID + Liveness	$< .001^*$
	vs. AI	$< .001^*$
	vs. Email	$< .001^*$
Gov-ID (NR)	vs. Gov-ID (SR)	1.00
	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	1.00
	vs. AI	0.82
	vs. Email	0.35
Gov-ID (SR)	vs. Gov-ID (CR)	1.00
	vs. Gov-ID + Liveness	0.64
	vs. AI	1.00
	vs. Email	1.00
Gov-ID (CR)	vs. Gov-ID + Liveness	0.49
	vs. AI	1.00
	vs. Email	1.00
Gov-ID + Liveness	vs. AI	0.29
	vs. Email	0.06
AI	vs. Email	1.00

A.3 Regression Models

Table 7: Ordinal Logistic Regression Results: Predictors of **Comfort** with Age Verification. Comfort was treated as an ordinal outcome on a five-point Likert scale (higher values indicate greater comfort). Age was coded 1-5 (higher values indicate older age brackets, with ages 51+ collapsed into 5) and education was coded 1-6 (higher values indicate higher educational attainment, with "Prefer not to respond" coded as missing). State law enactment was coded as a binary variable (0 = no enacted law, 1 = enacted law). The reported sample size ($N = 861$) excludes participants who selected "Prefer not to respond" for education. **Model fit**: log-likelihood $= -1143.40$, likelihood ratio: $\chi^2(9) = 297.70$, $p < 0.001$, $N = 861$. * indicates $p < 0.05$.

Predictor	**Coef**	**Odds Ratio**	**95% CI**	**z-value**	**p-value**
Age Group	0.21	1.23	[1.08, 1.40]	3.03	0.002*
Education Level	-0.07	0.94	[0.84, 1.04]	-1.20	0.231
State Law Enacted	0.05	1.05	[0.82, 1.37]	0.40	0.687
Verification Method (ref: Checkbox)					
Gov-ID (No Reassurance)	-2.70	0.07	[0.04, 0.11]	-10.54	<0.001*
Gov-ID (Simple Reassurance)	-2.89	0.06	[0.03, 0.09]	-11.37	<0.001*
Gov-ID (Compound Reassurance)	-2.89	0.06	[0.03, 0.09]	-11.70	<0.001*
Gov-ID + Liveness Check	-3.16	0.04	[0.03, 0.07]	-12.19	<0.001*
AI Facial Age Estimation	-2.54	0.08	[0.05, 0.12]	-11.44	<0.001*
Email Age Estimation	-1.54	0.22	[0.15, 0.32]	-7.63	<0.001*

Table 8: Ordinal Logistic Regression Results: Prior Experience and **Comfort** with Age Verification. Comfort was treated as an ordinal outcome on a five-point Likert scale (higher values indicate greater comfort). Past Age Verification Count is the total number of age verification methods participants reported having previously completed across 10 access contexts. R-rated Media Viewing Frequency (coded 0-4) and Prior Checkbox, Face-Based, and ID-Based Age Verification (coded 0-3) use higher number to indicate higher frequency and prior exposure, with "I am not sure" treated as missing. The reported sample size ($N = 803$) excludes participants who selected "I am not sure" for prior verification frequency. **Model fit:** log-likelihood $= -1045.30$, likelihood ratio: $\chi^2(11) = 307.51$, $p < 0.001$, $N = 803$. * indicates $p < 0.05$.

Predictor	**Coef**	**Odds Ratio**	**95% CI**	**z-value**	**p-value**
Past Age Verification Count	0.04	1.04	[0.97, 1.12]	1.03	0.303
R-rated Media Viewing Frequency	-0.05	0.95	[0.83, 1.10]	-0.65	0.518
Prior Checkbox Age Verification	-0.13	0.88	[0.74, 1.05]	-1.40	0.163
Prior Face-Based Age Verification	0.07	1.07	[0.85, 1.36]	0.57	0.566
Prior ID-Based Age Verification	0.45	1.57	[1.28, 1.92]	4.36	<0.001*
Verification Method (ref: Checkbox)					
Gov-ID (No Reassurance)	-2.80	0.06	[0.04, 0.10]	-10.42	<0.001*
Gov-ID (Simple Reassurance)	-3.08	0.05	[0.03, 0.08]	-11.51	<0.001*
Gov-ID (Compound Reassurance)	-3.04	0.05	[0.03, 0.08]	-11.83	<0.001*
Gov-ID + Liveness Check	-3.25	0.04	[0.02, 0.07]	-12.06	<0.001*
AI Facial Age Estimation	-2.78	0.06	[0.04, 0.10]	-11.74	<0.001*
Email Age Estimation	-1.64	0.19	[0.13, 0.29]	-7.73	<0.001*

Table 9: Ordinal Logistic Regression Results: Predictors of **Perceived Risk** with Age Verification. Risk was treated as an ordinal outcome on a five-point Likert scale (higher values indicate greater perceived risk). Past Age Verification Count is the total number of age verification methods participants reported having previously completed across 10 access contexts. Age was treated as a continuous predictor and coded 1-5 (higher values indicate older age brackets, with ages 51+ collapsed into 5). **Model fit**: log-likelihood $= -1226.10$, likelihood ratio: $\chi^2(8) = 280.86$, $p < 0.001$, $N = 881$. * indicates $p < 0.05$.

Predictor	**Coef**	**Odds Ratio**	**95% CI**	**z-value**	**p-value**
Past Age Verification Count	-0.02	0.98	[0.91, 1.04]	-0.73	0.465
Age Group	-0.09	0.92	[0.82, 1.03]	-1.41	0.159
Verification Method (ref: Checkbox)					
Gov-ID (No Reassurance)	2.67	14.48	[8.80, 23.82]	10.52	<0.001*
Gov-ID (Simple Reassurance)	2.61	13.54	[8.40, 21.84]	10.68	<0.001*
Gov-ID (Compound Reassurance)	2.65	14.18	[8.87, 22.68]	11.06	<0.001*
Gov-ID + Liveness Check	2.80	16.40	[10.07, 26.71]	11.24	<0.001*
AI Facial Age Estimation	2.05	7.75	[5.09, 11.79]	9.56	<0.001*
Email Age Estimation	0.77	2.15	[1.48, 3.13]	4.00	<0.001*

Table 10: Ordinal Logistic Regression Results: Predictors of **Perceived Convenience** of Age Verification. Convenience was treated as an ordinal outcome on a five-point Likert scale and was only assessed among those who completed age verification (higher values indicate higher perceived convenience). Reference categories are Checkbox (verification method) and Desktop (device type). Age Group was coded 1-5 (higher values indicate older age brackets, with ages 51+ collapsed into 5). ID Access Ease was coded 1-4 (higher values indicate easier access to government-issued ID), with "I'm not sure" coded as missing. The reported sample size ($N = 645$) excludes participants who selected "I'm not sure" in the ID Access Ease question. **Model fit**: log-likelihood $= -776.35$, likelihood ratio: $\chi^2(10) = 149.73$, $p < 0.001$, $N = 645$. * indicates $p < 0.05$.

Predictor	**Coef**	**Odds Ratio**	**95% CI**	**z-value**	**p-value**
Age Group	0.01	1.01	[0.86, 1.18]	0.07	0.943
ID Access Ease	0.64	1.90	[1.47, 2.47]	4.85	<0.001*
Device: Mobile	0.22	1.25	[0.88, 1.78]	1.24	0.216
Device: Tablet	1.08	2.95	[0.26, 33.54]	0.87	0.382
Verification Method (ref: Checkbox)					
Gov-ID (No Reassurance)	-1.92	0.15	[0.08, 0.27]	-6.14	<0.001*
Gov-ID (Simple Reassurance)	-1.98	0.14	[0.08, 0.25]	-6.69	<0.001*
Gov-ID (Compound Reassurance)	-2.52	0.08	[0.04, 0.15]	-8.25	<0.001*
Gov-ID + Liveness Check	-2.93	0.05	[0.03, 0.11]	-8.51	<0.001*
AI Facial Age Estimation	-0.93	0.39	[0.24, 0.64]	-3.73	<0.001*
Email-Based	-0.79	0.46	[0.29, 0.72]	-3.34	0.001*

Table 11: Logistic Regression: Association Between Perceptions and Age Verification Completion. Reference category for verification method is Gov-ID (No Reassurance). Comfort, Risk, and Effectiveness are treated as continuous independent variables (coded 1-5) where higher values indicate higher comfort, higher perceived risk, and higher perceived effectiveness, respectively. We used Gov-ID (NR) as the reference for this model to avoid quasi-complete separation in the logistic regression (all checkbox participants in the survey data completed age verification, and we omitted checkbox coefficients due to this perfect separation). Full results including the verification method coefficients are in [23]. **Model fit**: null deviance $= 983.63$ (df $= 880$), residual deviance $= 622.88$ (df $= 871$), likelihood ratio: $\chi^2(9) = 360.75$, $p < 0.001$, $N = 881$. * indicates $p < 0.05$.

Predictor	**Coef**	**Odds Ratio**	**95% CI**	**z-value**	**p-value**
Intercept	2.27	9.66	[2.05, 48.23]	2.82	0.005*
Comfort	0.35	1.42	[1.12, 1.82]	2.89	0.004*
Risk	-0.71	0.49	[0.37, 0.64]	-5.07	<0.001*
Effectiveness	0.05	1.05	[0.91, 1.22]	0.66	0.511