Measuring User Responses to Age Verification Architectures: Evidence from a Deceptive Online Experiment

Yanzi Lin, Vivianna Lieu, Cheng Zhang, Weiqian Zhang, Lorrie Faith Cranor, Sarah Scheffler

Carnegie Mellon University

Abstract

Following the U.S. Supreme Court's decision in Free Speech Coalition v. Paxton (2025), which established that age verification systems must be "adequately tailored," understanding user behavior has become legally relevant for system design. This preliminary study empirically examines how different age verification methods affect user behavior through a deceptive online experiment framed as usability testing for a mock gambling website. Participants (n=99 U.S. residents) were randomly assigned to six verification conditions, including simple checkbox self-declaration, government-issued ID upload, and AI-based facial age estimation. Results show stark differences in user responses: checkbox verification achieved 95.2% completion rates, while government ID methods drove up to 60.5% of users to return their study without finishing. We also tested the effects of privacy disclosures on completion rates. These had mixed effects, with detailed data handling information both increasing completion rates and polarizing user comfort levels. In a survey accompanying the empirical study, participants expressed significant privacy concerns about document-based methods, citing fears of identity theft and data misuse. These findings provide empirical evidence that can be applied to the U.S. Constitutional requirement for "adequate tailoring" of age verification systems, as well as policy analysis and technical design of age verification more broadly. We outline plans for expanded research using R-rated movie content to examine these effects at larger scale.

1 Introduction

The U.S. Supreme Court's decision in Free Speech Coalition v. Paxton (2025) established that age verification systems must be "adequately tailored" to avoid undue burdens on adults' Constitutional speech rights while effectively preventing minors' access to sexually explicit content that is obscene from a child's perspective [6]. This intermediate scrutiny standard makes user behavior constitutionally relevant, as the Court recognized that verification requirements may deter users who "fear for their reputations should the operator, advertently or inadvertently, disclose" information about their viewing habits [6].

Current age verification approaches represent distinct technical architectures with different implications for user privacy and compliance. These range from simple checkbox self-declaration to more complex systems involving government-issued identification documents, third-party verification services, and emerging biometric technologies like AI-based facial age estimation. These approaches create tradeoffs among security, usability, privacy, and other concerns that technical standards address through requirements for data minimization and proportional risk assessment [8].

While surveys consistently show public support for age verification that addresses privacy concerns [4, 16, 23], few studies have examined actual user behavior in realistic scenarios. For example, Ofcom found in 2022 that UK adults broadly support requiring age checks for adult sites, but express strong reservations about platforms collecting or misusing their personal data [16]. Similarly, surveys by the eSafety Commissioner in Australia report that while 78% of adults support age checks for pornography, many users lack awareness of how these systems operate and express concerns about who manages their data [4]. Research from the Family Online Safety Institute also shows that while biometric age estimation raises privacy concerns, a majority of parents and adolescents in the U.S. and U.K. are open to its use when framed as secure and privacy-preserving [5].

However, these surveys measure stated preferences rather than actual behavior when users encounter age ver-

IAB/W3C Workshop on Age-Based Restrictions on Content Access. October 7–9, 2025, London, UK.

ification requirements. Barry et al. [3] analyzed alcohol-brand websites and found that self-assertion mechanisms (e.g., entering a birthdate) were easily bypassed, with most sites allowing unlimited retries when users initially failed age verification. Stewart et al. [19] similarly report that minors face few practical obstacles in circumventing online age restrictions, often doing so through repeated attempts or by providing fabricated information. To our knowledge, no controlled experiments to date have directly compared how users respond to multiple age verification methods in the moment of decision-making.

This study addresses that gap through a deceptive online experiment measuring how users respond to six different age verification approaches. Using a mock gambling site with third-party age verification, we observe completion rates, study return behavior, and use of an alternative verification method to understand how different age verification mechanisms affect user compliance. Our findings provide empirical evidence relevant to both the legal "adequately tailored" standard and technical system design decisions.

We report findings from our pilot study with 99 U.S.-based participants and outline design modifications for a larger-scale study. Specifically, we investigate the following questions in our pilot study:

(1) How do different age verification methods affect users' decisions to access an online gambling site?
(2) How do disclosures of data handling practices affect those decisions?
(3) What are users' attitudes toward different age verification mechanisms?

2 Pilot Study Methods

2.1 Study Protocols

We conducted a deceptive, 99-participant, two-part online study, framed as a usability test for the mock gambling website "Bet Sierra" and using a mock age verification service "AgeGuardian." This setup simulated real-world online age verification, with gambling selected as the context because it represents a common use case for age verification while avoiding the heightened sensitivities associated with sexual content.

We recruited individuals aged 21 and older, fluent in English, and residing in the U.S., via the online crowdworker platform Prolific. Participants received $1.25 for completing the experimental part and an additional $1.25 for the follow-up survey (consistent with receiving approximately $15/hour).

In the experiment, participants were redirected to *betsierra.com*, where they were informed through a pop-up that they must verify their age before proceeding. They were then sent to *ageguardian.net*, a mock verification site presented as a third-party service (however in reality, both websites were controlled by the researchers). Participants were randomly assigned to one of six age verification conditions based on common methods used in recent legislation and commercial platforms (see Table 1). These conditions were selected to represent the range of age verification approaches currently or recently implemented.

Although the mock age verification website presented a prompt that appeared to collect participants' relevant data needed for the age verification condition (e.g., a picture of their ID), in reality the mock website did not actually collect or transmit any of this information to the researchers or to any other party.

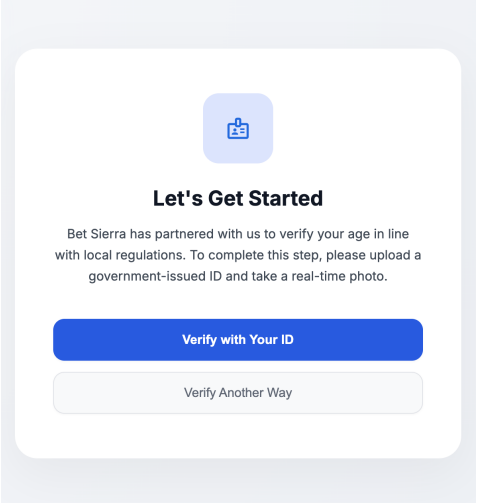

Figure 1: Example AgeGuardian interface for the Gov-ID + Real-Time Photo condition.

On the AgeGuardian site, participants could either complete their assigned verification method (example interfaces in Figure 1 and Figure 2) or choose an "alternative method," which they were told involved entering their Prolific ID and waiting 24 hours for verification by the research team. This alternative method was designed to capture participants who refused to complete their assigned method but remained interested in accessing the gambling site through different means despite facing more inconvenience (i.e., a 24-hour delay in access and payment, and the apparent need to manually follow up).

In addition to testing four primary age verification methods, we also tested three privacy reassurance variations. These tested whether specific disclosures about data handling practices (i.e., data use purposes, data retention, and data sharing policies) would affect users' age verification behavior and attitudes, independent of the underlying technical system.

Table 1: Verification Conditions and Corresponding Prompts

Assurance Method	**Opening Prompt**
Checkbox Age Gate	"I certify that I am the minimum legal age to gamble in my jurisdiction."
Government-Issued ID with No Reassurance	"Please upload a photo of your government-issued ID."
Government-Issued ID with Simple Reassurance	(+) "Your data will only be used for age verification."
Government-Issued ID with Compound Reassurance	(+) "Your data will be securely stored and deleted immediately after the verification process is complete. In compliance with local laws, we will not share your data with 3rd party data brokers."
Government-Issued ID and Real-Time Photo	(+) "and take a real-time photo for verification."
AI Facial Age Estimation	"Please upload a photo of yourself to verify your age using AI analysis."

All "Government-Issued ID" conditions begin with "Please upload a photo of your government-issued ID" followed by statement after +.

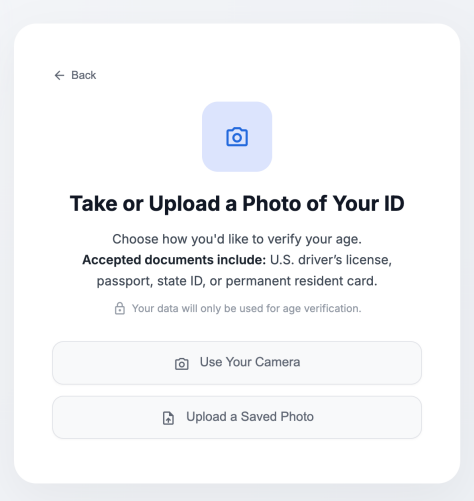

Figure 2: Example photo ID upload interface for the Gov-ID with Simple Reassurance condition.

After completing the verification process (or choosing to do the alternative method), participants were redirected to a debriefing page where the study's true purpose was revealed. Following the debrief, participants could choose to withdraw their data without losing compensation or proceed to the follow-up survey. The survey collected information about their attitudes toward age assurance, comfort with the assigned method, factors influencing their verification decisions, prior experiences with online age verification, and demographic characteristics.

2.2 Ethical Considerations

The pilot study was approved by the Carnegie Mellon University Institutional Review Board.

Our pilot study involved the use of deception. During the consent process, we informed participants that some aspects of the study's purpose or procedures could not be fully disclosed upfront to avoid influencing their responses. All participants provided informed consent before proceeding to the main study. After completing the age verification process, participants underwent a thorough debrief where we disclosed that the true focus was their interactions with AgeGuardian, not website usability testing. Participants were informed that no actual verification data (e.g., selfies or ID photos) was collected, stored, or transmitted from their devices and that the upload process was entirely simulated. Those assigned to the AI age estimation condition were told no AI model was used to estimate their age. We also informed participants that we had collected web activity data including their clicks and completion times during the verification process. At this point, participants were given the option to withdraw their data if they felt uncomfortable with the deception or initial incomplete disclosure.

The AgeGuardian verification website was developed by the research team and programmed to ensure no actual data transfer occurred. After data collection was complete, the research team removed all identifiable information and securely stored the anonymized data (survey responses and website interactions like clicks and timing) in a shared Google Drive folder with access restricted to research team members.

3 Pilot Results

As a preliminary study, we recruited a gender-balanced sample of 99 U.S.-based participants, with at least 15 in each of six conditions. Participants were predominantly White (72.7%), with 20.2% Black, 5.1% Asian and 3.0% Other Race. 74.7% held at least a Bachelor's degree. 70.7% had prior online gambling experience, with 74.3% of those visiting gambling sites at least monthly. The average completion time was 12.5 minutes. One participant opted out of the survey but chose not to withdraw their data from the experiment. An additional 68 participants "returned" the study on Prolific after being asked to age verify, leaving the study without completing it. Because returning the study, in many cases, likely signals users' behavior toward age verification, we include it in the study's results (see Table 2 and Figure 3). In our ongo-

ing follow-up to this study, we are designing the study to channel people who might have "returned" the study out of discomfort with age verification to take an action more akin to the "alternative method" rather than study withdrawal.

Table 2: Number of Approved and Returned Participants by Condition

Condition	Approved	Returned
Checkbox	20	1
Gov-ID with No Reassurance	15	23
Gov-ID with Simple Reassurance	15	18
Gov-ID with Compound Reassurance	15	9
Gov-ID + Real-Time Photo	18	8
AI Age Estimation	16	6

Note: Age assurance methods are abbreviated. "Approved" refers to the participants who consented to retain their data and were approved for payment on Prolific. "Returned" refers to participants who returned their study on Prolific without finishing.

RQ1: Among all age assurance methods, the Checkbox condition–where participants simply checked a box to confirm their age–had the highest completion rate (95.2%) with only 4.8% returning the study and 0% using the alternative method. This condition also showed the highest reported comfort, with 71.4% of participants feeling at least "Somewhat Comfortable" with this method.

In contrast, the Government-Issued ID with with No Reassurance condition showed the highest return rate (60.5%), while only 3.0% of participants in the Simple Reassurance condition completed the assigned verification process. The Government-Issued ID with Real-Time Photo condition had the highest proportion of participants verifying via the alternative method. However, over half of those reported lacking access to a camera. These findings suggest that as verification methods become more invasive or technically demanding, users become increasingly likely to abandon the verification process entirely or seek less intrusive alternatives. Complete results on completion rates and reported comfort are in Figures 3 and 4.

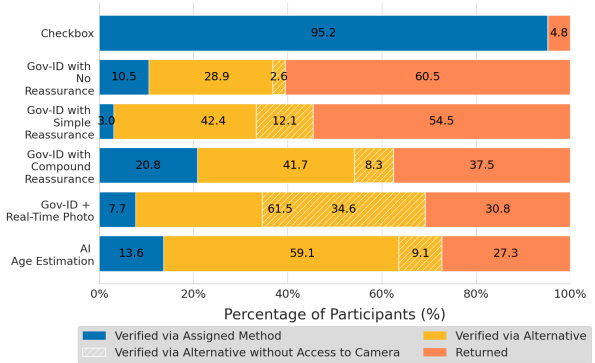

Figure 3: Percentage of people within each condition that verified via the assigned method, the alternative method, or returned.

RQ2: Disclosures about data handling practices had mixed effects on users' access decisions. The Compound Reassurance condition, which included information about both the purpose of data collection and how data would be handled, resulted in the highest completion rate among the Government-Issued ID conditions. Interestingly, this condition also produced the most polarized comfort responses, with the highest percentages of participants reporting being "Very Comfortable" (12.5%) and "Very Uncomfortable" (12.5%). However, discomfort may be underrepresented in the No Reassurance and Simple Reassurance conditions, as participants who felt highly uncomfortable in these conditions may have returned the study before providing comfort ratings in the survey.

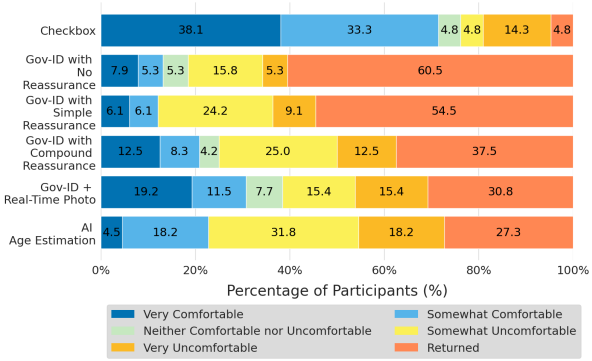

Figure 4: Self-reported comfort ratings across verification methods.

RQ3: Many participants expressed privacy and security concerns regarding online age assurance, particularly with methods requiring personal identification documents. Many were reluctant to upload government-issued IDs due to fears of identity theft and data security risks. Trust was a crucial factor influencing attitudes, with a participant explicitly stating "I won't be sharing any personal information or images with companies that I don't trust." This trust concern was compounded by confusion about the multi-party architecture: participants were uncertain whether the gambling site (Bet Sierra) or the third-party verification service (AgeGuardian) was collecting their data. This highlights a key challenge in hybrid verification architectures, where multiple organizational boundaries can obscure user understanding of data flows and accountability. Participants in the Checkbox condition acknowledged the trade-off between ease of access and verification accuracy: "Not requiring additional verification beyond my confirmation made it seamless for me, but I have concerns with its ability to catch people who are under 21 who will lie." Many specifically objected to AI-based verification, describing it as "stressful," questioning what other data might be collected beyond age estimation, and expressing concerns that "AI can glitch and ruin things for a lot of people."

3.1 Limitations

We identify several limitations with our pilot study. First, the small sample size (n=99) limits our ability to draw quantitative conclusions. Second, we did not restrict recruitment to individuals with prior online gambling experience. Participants without online gambling experience were likely less familiar with age verification requirements and less motivated to access the gambling site, which may have reduced completion rates and comfort levels. Third, our deceptive design created study legitimacy concerns. One participant noted in the post-experiment survey their confusion about the simulated ID upload, as Prolific explicitly prohibits collecting sensitive personal information that risks de-anonymizing users. Although no actual age verification data were collected, the use of a deceptive study design created the perception of a violation of Prolific's platform policies, which may have made the study appear illegitimate and led to more participants "returning" their study on Prolific.

We address these limitations in the following section.

4 Next Steps

Our preliminary study provided initial insights into user age verification behavior in an online gambling context and informed our approach for the full-scale study. We are implementing several design modifications based on our pilot findings.

Use Case: While our pilot study focused on gambling websites, the Supreme Court's decision in FSC v. Paxton creates uncertainty about how age verification requirements might expand beyond pornographic content. The Court explicitly noted that H.B. 1181 itself "cannot conceivably be read to cover, say, a PG-13- or R-rated movie," [6] but this distinction may not hold for other broader age verification mandates. R-rated movies occupy a unique middle ground—they contain sexual content that approaches the threshold for age verification and already require ID checks or parental accompaniment in theaters, yet have no equivalent online age verification requirements. This makes them a valuable test case for studying user responses to potential expanded age verification mandates. We will advertise the study as seeking participants to view clips from R-rated romantic films and subsequently complete a survey to provide feedback. This framing offers an authentic research context while masking our focus on age verification behaviors.

Experimental Conditions: Building on our six pilot conditions, we will add Email-Based Age Estimation as a seventh condition to test "transactional data" verification methods referenced in Arizona H.B. 2112 [2] and and other U.S. age verification laws that define transactional data verification in similar ways [1, 2, 7, 9–15, 17, 18, 20–22]. Arizona H.B. 2112 defines transactional data as "a sequence of information that documents an exchange, agreement, or transfer between an individual, commercial entity or third-party entity" and includes "records from mortgage, education and employment entities."

For this condition, we will model our approach on that used by Yoti, a UK-based digital identity company that offers email-based age estimation. Participants will be informed that their email addresses will be checked against third-party databases to analyze the email's registration source, associated employer information, and linked financial activities to estimate age. This addition will provide insights into user responses to transactional data verification, as practical implementation of this method remains unstandardized despite its legal recognition.

Sample Size: We plan to increase our sample size based on power analysis to provide sufficient statistical power for comparing user responses across different age verification mechanisms.

Recruitment Platforms: Our pilot study used Prolific, but the use of a deceptive design raised concerns about study legitimacy, as discussed in Section 3.1. To address this, we contacted Prolific via email to explain the purpose of our study and clarify that the data collection was only simulated. We also requested approval to display a yellow warning to future participants to show that the study is pre-verified by the platform to collect personal information. After the email exchange, Prolific decided that the design could cause participant distress and damage platform trust, regardless of whether data were actually collected. Therefore, for our full study, we will recruit through online advertising platforms such as Facebook Ads to reach participants beyond crowdworkers. These platforms allow targeted recruitment based on interests in romantic media or adult content, which will enable us to study users who would naturally encounter age verification requirements for R-rated content.

Survey Design: Our pilot survey focused primarily on user comfort with different verification methods. The expanded survey will include additional constructs to provide a more comprehensive understanding of user attitudes: perceived effectiveness of different methods at blocking underage users, perceived privacy risks associated with uploading personal documents, ease of accessing required verification materials, and user preferences for non-self-assertion verification methods when accessing R-rated content online.

We anticipate deploying the full study in Fall 2025, though the design may continue to evolve as we incorporate feedback from the research community and respond to ongoing developments in age verification policy and technology.

References

[1] Arkansas senate bill 66 (2023), act 612. https://arkleg.state.ar.us/Home/FTPDocument?path=%2FACTS%2F2023R%2FPublic%2FACT612.pdf, 2023.

[2] Arizona house bill 2112 (2025), chapter 193, laws of 2025. https://www.azleg.gov/legtext/57leg/1R/laws/0193.pdf, 2025. 57th Legislature, 1st Regular Session.

[3] Adam E. Barry, Allison M. Bates, Oluwafemi Olusanya, Carli E. Vinal, Amanda Martin, Jessica E. Peoples, Jasmary Montano, and Elizabeth H. Chaney. An Evaluation of Alcohol Brand Website Age Gates: Effectiveness and Implications. *Alcoholism Treatment Quarterly*, 39(3):313–327, 2021.

[4] eSafety Commissioner. Public perceptions of age verification for limiting access to pornography. https://www.esafety.gov.au/sites/default/files/2021-10/Public%20perceptions%20of%20age%20verification%20fact%20sheet.pdf?v=1754497181520, 2021.

[5] Family Online Safety Institute. Making Sense of Age Assurance: Cross-national Survey Results from the US, UK, and France. https://www.fosi.org/policy-research/making-sense-of-age-assurance, 2022.

[6] Free Speech Coalition v. Paxton, 603 U.S. __ (2025). https://www.supremecourt.gov/opinions/24pdf/23-1122_3e04.pdf, 2025. U.S. Supreme Court Decision, No. 23-1122, decided June 26, 2025.

[7] Idaho house bill 498 (2024). https://legislature.idaho.gov/wp-content/uploads/sessioninfo/2024/legislation/H0498.pdf, 2024.

[8] IEEE Standard 2089.1-2024: IEEE Standard for Online Age Verification. https://ieeexplore.ieee.org/document/10542699, 2024. IEEE Std 2089.1-2024.

[9] Indiana senate enrolled act no. 17 (2024). https://iga.in.gov/pdf-documents/123/2024/senate/bills/SB0017/SB0017.05.ENRH.pdf, 2024.

[10] Kentucky house bill 278 (2024), chapter 106. https://apps.legislature.ky.gov/law/acts/24RS/documents/0106.pdf, 2024.

[11] Louisiana house bill 142 (2023). https://www.legis.la.gov/Legis/ViewDocument.aspx?d=1289498, 2023.

[12] Mississippi senate bill 2346 (2023). https://billstatus.ls.state.ms.us/2023/pdf/history/SB/SB2346.xml, 2023.

[13] Montana senate bill 544 (2023). https://archive.legmt.gov/bills/2023/billhtml/SB0544.htm, 2023.

[14] North dakota house bill 1561 (2025). https://ndlegis.gov/assembly/69-2025/regular/documents/25-0968-03000.pdf, 2025.

[15] Nebraska legislative bill 1092 (2024). https://nebraskalegislature.gov/FloorDocs/108/PDF/Slip/LB1092.pdf, 2024.

[16] Ofcom. Adult users' attitudes to age-verification on adult sites. https://www.ofcom.org.uk/__data/assets/pdf_file/0019/239103/adult-attitudes-age-verification.pdf, 2022.

[17] Oklahoma senate bill 1959 (2024). https://www.oklegislature.gov/cf_pdf/2023-24%20ENR/SB/SB1959%20ENR.PDF, 2024.

[18] South carolina bill 3424 (2023–2024). https://www.scstatehouse.gov/sess125_2023-2024/bills/3424.htm, 2023.

[19] Hannah Stewart, Roslyn Campbell, Jane Carr, and Nathan Grayson. Feasibility and acceptability of age verification technologies: Insights from qualitative research with parents and adolescents. Journal of Child and Family Studies (in press), 2024.

[20] Tennessee senate bill 179 (2024), public chapter 1021. https://publications.tnsosfiles.com/acts/113/pub/pc1021.pdf, 2024.

[21] Texas house bill 1181 (2023), 88th legislature, regular session. https://capitol.texas.gov/tlodocs/88R/billtext/pdf/HB01181F.pdf#navpanes=0, 2023.

[22] Utah senate bill 287 (2023). https://le.utah.gov/~2023/bills/static/SB0287.html, 2023.

[23] Paul J. Wright and Debby Herbenick. U.S. Attitudes Toward Age-Verification for Online Pornography. Journal of Sex Research (forthcoming), 2025.