User (Non-)Compliance with Age Verification: Preliminary Evidence from a Deceptive Web Experiment

Yanzi Lin, Cheng Zhang, Madelyne Xiao, Lorrie Faith Cranor, Sarah Scheffler[1]
Carnegie Mellon University
CyLab Security and Privacy Institute
January 2026

Online age-verification prompts increasingly common across the Internet, especially on social media and adult-content websites. As of January 2026, twenty-five U.S. states have active laws that require some websites to perform "strong" age verification on users to ensure that all visitors to sites containing "material harmful to minors" are over-age – and eleven more states have a bill announced or introduced. For the affected websites, these new laws mean that the previous common approach of self-attesting one's age by checking a box is no longer sufficient. Instead, users must verify their age by using IDs, AI facial analysis, or other "commercially reasonable" options.

However, the new approaches to age verification come with costs. Users may find them privacy-invasive, insecure, or inconvenient, and some users may even turn away from a website entirely if prompted with one of these new methods. But *how many* users will find these concerns overwhelming? Do users prefer one method or another? Do they believe the methods will be effective? These questions are of practical importance to websites that must now use age verification or will be required to do so soon. They are also relevant to determining whether specific age verification requirements are "adequately tailored" to their goal, legally balancing the incidental costs of age verification against restrictions on accessing content.

Our team at Carnegie Mellon University ran a study to find the answers to these questions. Our study recruited 1,635 participants to watch romantic movie clips and take a short research survey to be entered into a raffle. Participants were randomly assigned to age verify using one of seven approaches (five age assurance methods, one of which had three variants). Completion results varied significantly by age assurance method. We highlight some of the most important results in the bullets below. See Section 1 for our study methods and Section 2 for more results.[2]

- **When prompted to age verify with a government-issued ID, only 22-28% of users completed the process.** Adding a liveness check (requiring users to blink, move their head, or smile to verify that the uploaded ID corresponds to the person completing verification) reduced completion to 17%. AI-based facial analysis had a 52% completion rate, and an email-based approach in which we stated we would look up financial documents based on email had an 86% completion rate. In comparison, a self-attestation "checkbox" approach had a 99% completion rate. More details in Section 2.1.
- **The majority of participants were at least somewhat uncomfortable with all non-checkbox forms of age assurance, even among those who completed age assurance in our experiment.** Among age assurance completers, 53-80% of users were at least somewhat uncomfortable with the approach; among age assurance non-completers, this rose to 83-93%. More details in Section 2.2.
- **Participants did not think any age assurance method was very effective.** Participants agreed the checkbox was ineffective (only 8.3% said this method was even somewhat effective), but all methods had at most 40.5% of participants say it was at least "somewhat effective." More details in Section 2.3.
- **Participants found most Gov-ID conditions to be significantly less convenient than other methods.** The checkbox, email verification, and AI facial analysis verification were perceived to be more convenient than ID-based alternatives; the checkbox was ranked as significantly more convenient than all alternatives. More details in Section 2.4.

[1]Corresponding author: Sarah Scheffler, sscheffl@cmu.edu, assistant professor at CMU CyLab
[2]These results are preliminary and await final peer review.

1 What we did

Our primary goal was to empirically measure what fraction of people would complete age assurance when presented with different forms of assurance – and why. To measure this, we conducted a two-part study:

- **Part 1: Deceptive Web Experiment.** The first part of the study measured how many users would complete various methods of age assurance in a real-world setting. We measured how many people completed age assurance for each of seven age assurance conditions.
- **Part 2: Follow-Up Survey.** The second part of the study involved a follow-up survey, where we asked more detailed questions about the user's attitudes and past behaviors regarding age assurance.

The full study flow, from the experiment to the survey, is shown in Figure 1.

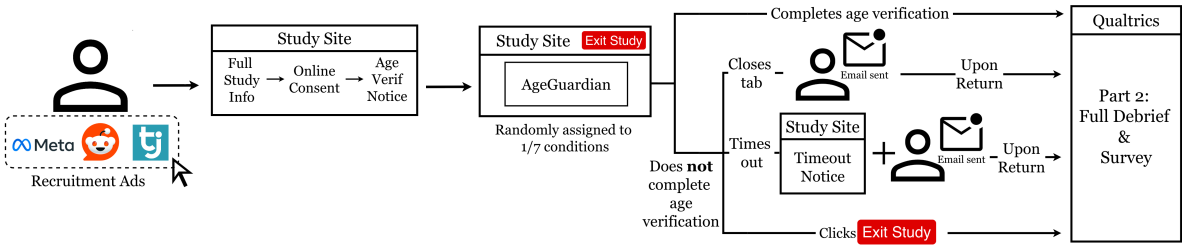

Figure 1. Study flow: Users were recruited via advertisements on multiple platforms. Those who clicked on the ads went to a website that presented study info, obtained their consent, and then presented them with an age verification notice (part 1). Those who age verified were directed to part 2 of the 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 do part 2.

1.1 Age assurance methods evaluated

We evaluated five age assurance methods: **(1) Checkbox Self-Attestation**, **(2) Government-Issued ID (Gov-ID)**, **(3) Gov-ID and Liveness Check** (included a "video selfie" liveness check in addition to an ID check), **(4) AI Facial Age Estimation**, and **(5) Email Age Estimation** (in which we informed users we would use their email to check the email's registration source, employer information, linked financial activities, educational records, and behavior patterns).

For the Gov-ID method, we included three variants: **(i) Simple Reassurance** (in which we additionally stated a purpose restriction on the data use), **(ii) Compound Reassurance** (in which we stated a purpose restriction and a promise on deletion of identifying information), and **(iii) No Reassurance** (in which we asked for an ID upload with no additional text).

In total we had seven conditions: three Gov-ID conditions plus the four other primary conditions. The text for all conditions can be found in Table 1, and screenshots of the Gov-ID interface are in Figure 4.

1.2 Study methods

The full study flow is shown in Figure 1. The study was approved by CMU's Institutional Review Board.

Recruitment. Study participants were recruited via ads posted on Meta, Reddit, and TrafficJunky (an advertising company that serves ads on adult entertainment websites) in November and December 2025. We recruited participants using ads asking users to "Watch Romantic Clips (18+) & Take A Short Research Survey" and be entered into a raffle to win up to $210. One of our ads is shown in Figure 2.

Part 1: Deceptive Web Experiment. In order to measure what people would actually do when asked to age verify, we designed the study so that study participants would think that the age verification was only incidental to the study. The study website interface is shown in Figure 3.

Table 1. Age Assurance Conditions and Corresponding Prompts

Assurance Method	Opening Prompt
Checkbox Self-Attestation	"I certify that I am 18 years of age or older and am legally permitted to view R-rated content."
Government-Issued ID, No Reassurance (NR)	"Please upload a government-issued ID. Accepted documents include: U.S. driver's license, passport, state ID, or permanent resident card."
Government-Issued ID, Simple Reassurance (SR)	(+) "Your data will be used only for age verification."
Government-Issued ID, Compound Reassurance (CR)	(+) "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	(+) "Please upload a government-issued ID and complete a liveness check to verify your identity matches your ID. "
AI Facial Age Estimation	"Please upload a photo of your face. We will use AI to estimate your age."
Email Age Estimation	"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."

All "Government-Issued ID" conditions begin with "Please upload a photo of your government-issued ID" followed by statement after (+). Screenshots of an AgeGuardian ID verification interface are in Figure 4.

Figure 2. One ad we used to recruit study participants

In this deceptive experiment, we informed participants that our research team at CMU had partnered with "AgeGuardian" to verify age before proceeding with the study. AgeGuardian screenshots are shown in Figure 4. In reality, AgeGuardian was a website built and controlled by the researchers for this study that did not actually collect photos or any personally-identifiable information.

Each participant was randomly assigned to one of seven age assurance conditions (see Table 1). Upon starting the study, the participant would be shown an embedded AgeGuardian interface for their age assurance method. For each condition, we measured how many people completed the the mock age assurance, versus how many people left by clicking "Exit Study," closing the tab, or timing out. We discuss results in Section 2.

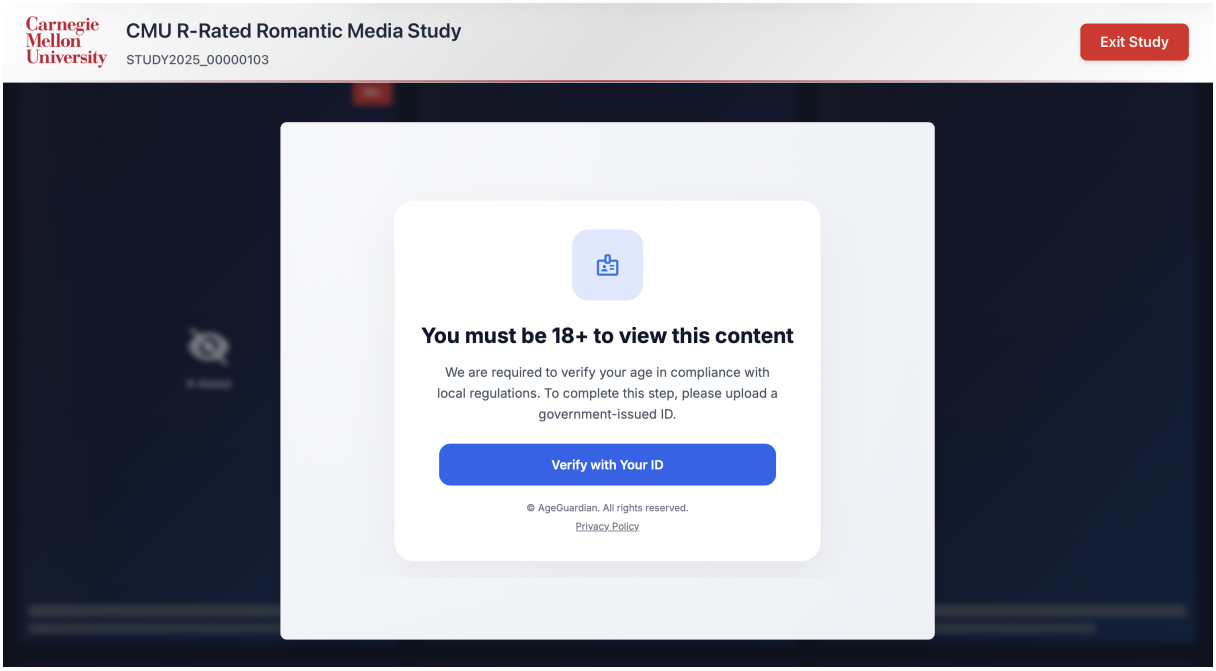

Figure 3. The study interface, shown here, was designed so that study participants would think that the age verification was only incidental to the study

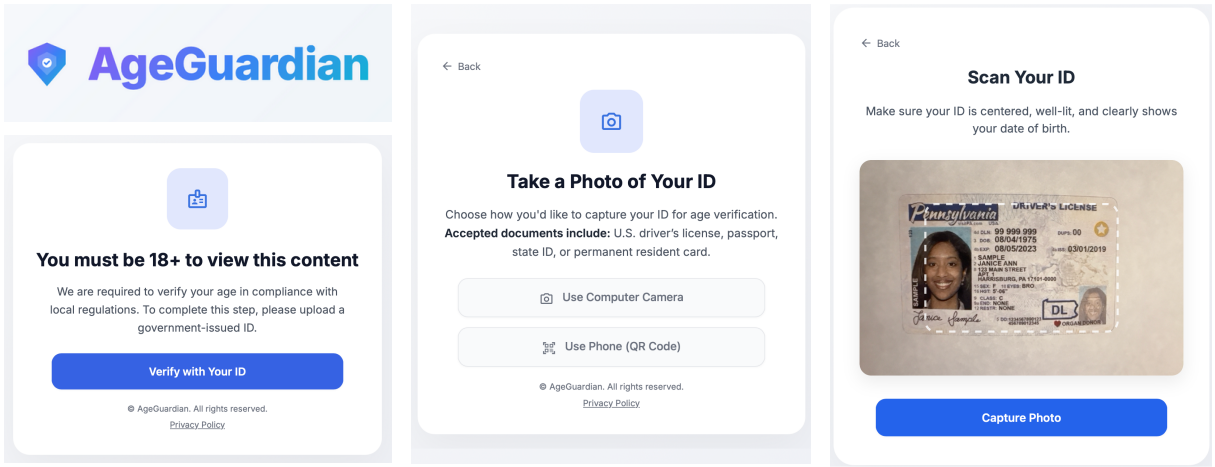

Figure 4. "AgeGuardian" interface shown to study participants in the Gov-ID with No Reassurance condition. The experiment measured what fraction of users completed this process for seven age assurance method conditions

Part 2: Survey. After age verification was either completed or rejected, participants were directed to a survey. To make sure we could collect responses from all participants, not just those who completed the age assurance process, we sent an email invitation to complete the survey to participants who quit in Part 1.

In the survey we asked participants questions about their attitudes toward age assurance, including whether they were comfortable with the age assurance method used, how convenient it is, how effective they think it is, and their past experience with it. We collected responses to attitude questions on five-point Likert scales along with free-text responses. We discuss results in Section 2.

2 What we found

We collected 1,635 valid responses to our web experiment and 881 valid responses to our follow-up survey after removing duplicates and suspected bots.

2.1 Age assurance completion rates vary significantly by method

As shown in Figure 5, age assurance completion rates varied significantly by verification method. Checkbox self-attestation achieved the highest completion (99.2%), while all government-ID based conditions showed substantially lower completion (22.2-27.3%), with the Gov-ID + liveness check condition having the lowest completion rate (17.9%). Email-based age estimation exhibited a significantly higher completion rate (86.0%) than all other methods except checkbox self-attestation. AI facial age estimation had the third highest completion rate (51.5%) among the age verification methods tested. Simple reassurance and compound reassurance did not have a significant impact on government-ID completion rates.

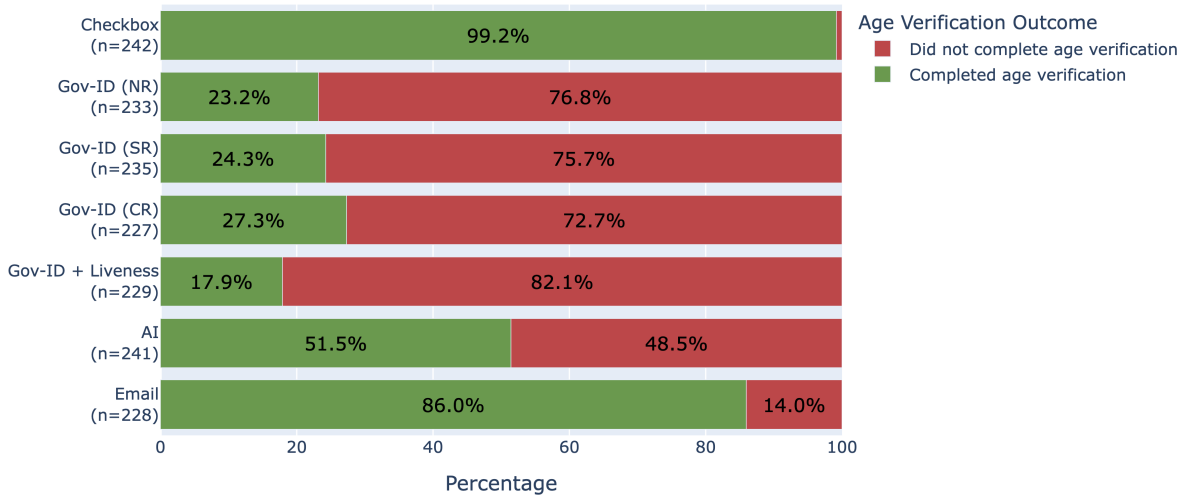

Figure 5. Completion rates for all age assurance methods. Completion rates varied statistically significantly across age assurance methods. In pairwise comparisons (Wilcoxon rank-sum tests, Holm-Bonforreroni adjusted p-values shown in Table 8), Checkbox completion was significantly higher than all other methods ($p < .001$). All Gov-ID conditions (including Gov-ID + Liveness) were significantly lower than AI and Email methods ($p < .001$). Completion of email was significantly higher than completion of AI ($p < .001$).

A chi-squared test of independence showed a strong association between verification method and completion outcome ($\chi^2 = 619.82$, df=6, $p < 0.001$) with a large effect size (Cramér's $V = 0.61$, 95% CI $[0.57, 1.00]$). Pairwise comparisons were statistically significant between Checkbox and all other conditions, between all Gov-ID conditions (including Gov-ID + Liveness) and all non-Gov-ID conditions, and between AI and email.

Pairwise comparisons revealed no statistically significant differences among the three government-issued ID reassurance conditions (No Reassurance, Simple Reassurance, and Compound Reassurance). In the survey, we asked participants whether reassurance statements affected how they felt about the age assurance process. Among participants who completed verification, most reported increased comfort after seeing the statements: 56.9% for Simple Reassurance condition and 48.2% for Compound Reassurance, while only 6.9% and 3.7% respectively reported decreased comfort. However, among participants who did not complete verification, reassurance statements had minimal impact. In both conditions, most reported no change in comfort (75.7% for Simple Reassurance, 56.2% for Compound Reassurance).

2.2 Users were at least somewhat uncomfortable with all non-checkbox methods

Comfort was measured by asking the question, "How comfortable were you with certifying your age through [method]?" and collecting responses on a five-point Likert scale. The results differed for those who completed vs didn't complete age verification, as shown in Figure 6.

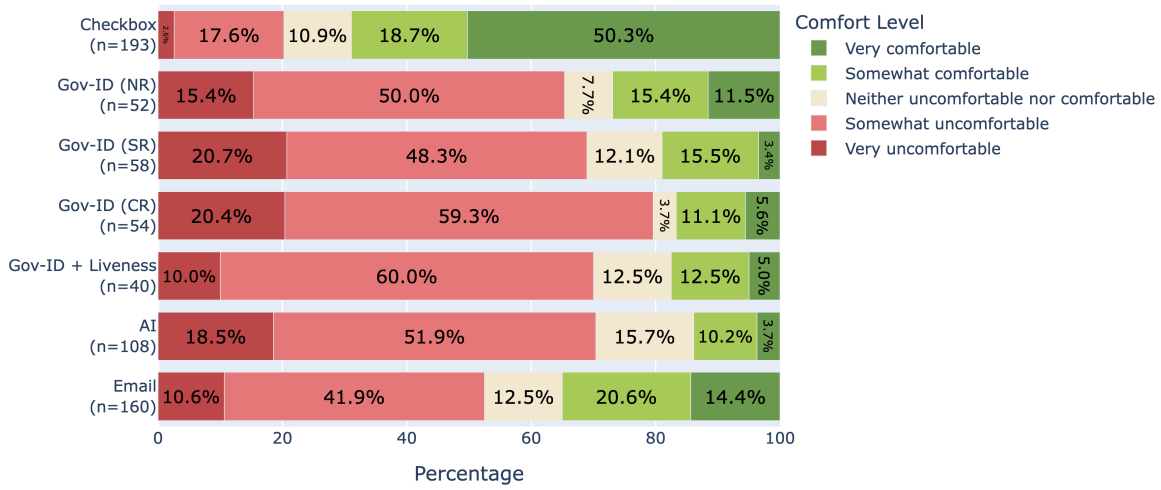

(a) Reported **comfort** among participants who **completed** age assurance

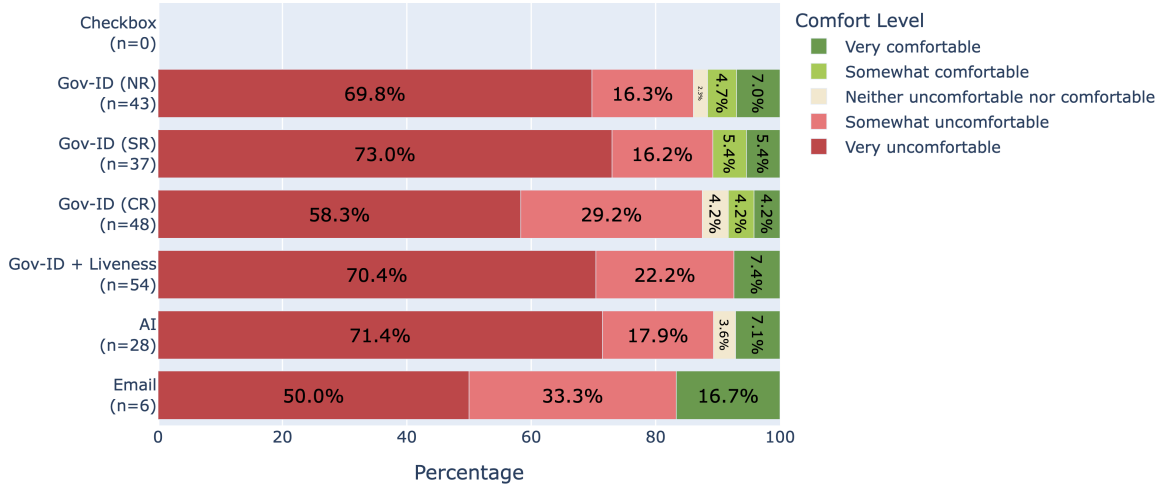

(b) Reported **comfort** among participants who **did not complete** age assurance

Figure 6. Reported **comfort** among participants who **completed** age assurance (top) and those who **did not complete** age assurance (bottom). Wilcoxon rank-sum tests with Holm-Bonferroni-adjusted p-values found statistically significant pairwise differences between Checkbox and all other conditions, and between email and all other conditions.

Overall, a majority of users were somewhat or very uncomfortable with all non-checkbox forms of age assurance – even among participants who completed the age verification. As shown in Figure 6, participants who completed age assurance tended to say they were "somewhat uncomfortable" while those who did not complete tended to say they were "very uncomfortable." Furthermore, as we note in Table 3, 68.5% of participants said their comfort was affected positively by knowing the information was used by

Carnegie Mellon University compared to some other entity, so these numbers are likely higher than they would be for an age verifier in the wild.

Statistically significant differences were found in comfort levels across conditions ($\chi^2 = 252.21$, $p < 0.001$). However, no significant differences emerged among the Gov-ID conditions (including Gov-ID with Liveness Check; all $p = 1.00$), or between AI and any Gov-ID condition ($p \geq 0.25$). Pairwise comparisons were performed using the Wilcoxon rank sum test with Holm-Bonferroni correction for multiple comparisons. Statistics are shown in Table 9 in Section A.2.

In open-ended responses, participants explained their comfort responses. Some of those who expressed being comfortable with age verification felt that it was less intrusive than other types of online data collection. One participant wrote, "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." However, most participants expressed concerns with providing any type of personally identifiable information to a verification service that could potentially use it for other purposes or share it with other companies. One participant wrote "[C]ategorically I will never be fully comfortable with age verification because it's a slippery slope to abusive data practices. PII is routinely linked publicly, typically with no remedy for those harmed." Another wrote, "It's an untrustworthy insecure medium to share those documents. Most id verification in person on [sic] requires a glance as verification rather than uploading photos. Given the proliferation of AI and deepfakes, I'd rather not upload an ID photo that could easily be used maliciously." Another noted, "[E]ach time you do it, it is through a different service. ... there may be 10's or 100's of these kinds of verification services, and at that point it only takes one to be a bad actor."

Some participants noted feeling more comfortable with the email method due a feeling that an email address was not as intrusive as a facial scan or drivers license. For example, one participant wrote, "[Email verification] identifies me on the internet but doesn't necessarily identify me in real life, I suppose." Another wrote "I don't love the idea that you can figure out my age based on my email in just a few seconds, but I do realize the importance of a quick and relatively less intrusive way of age verification." But others remained concerned about giving their email address to a website they weren't familiar with. one wrote, "I am weary [sic] about giving any data of mine (even just my email) to a company/website about which I know nothing."

2.3 Users perceived limited effectiveness for all methods

We measured effectiveness by asking participants, "How effective do you think this verification method is at preventing underage users from accessing R-rated media content online?" on a five-point Likert scale. Results by condition are shown in Figure 7.

Participants widely agreed that the checkbox was not effective (8.3% said at least somewhat effective), and was rated as statistically significantly less effective than all other methods (all $p < 0.001$). However, participants also found the other approaches to be only somewhat effective. Participants thought the Gov-ID + Liveness ("video selfie") approach was most effective (40.5% said at least somewhat effective), followed by Gov-ID in general (38.9%). Other approaches fell between this range (see Figure 7). No significant differences in perceived effectiveness were found among any of the Gov-ID conditions (all $p \geq 0.49$).

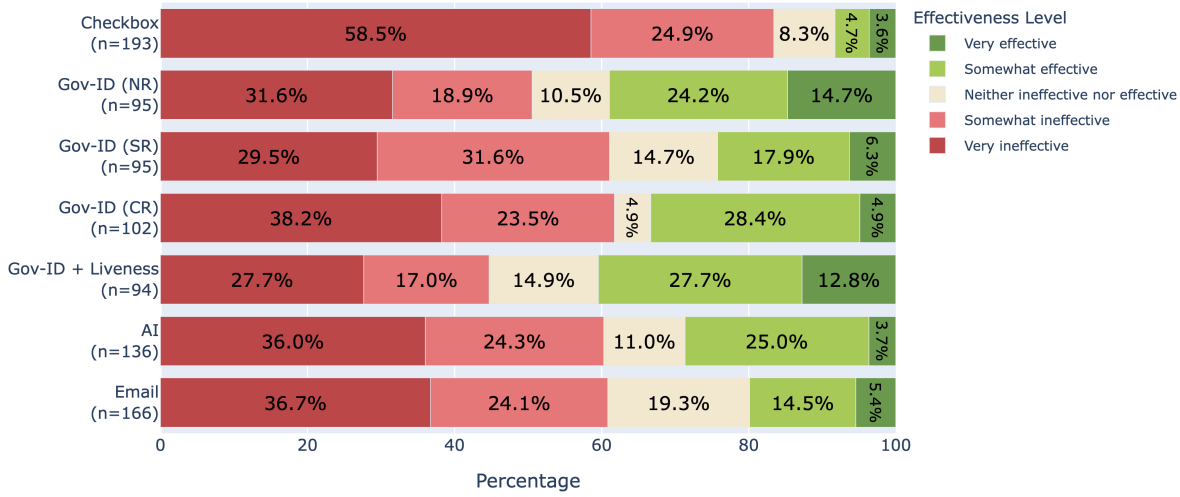

Figure 7. Perceived **effectiveness** for all age assurance methods. In pairwise comparisons (Wilcoxon rank-sum tests, Holm-Bonferreoni adjusted p-values shown in Table 10, Checkbox was perceived as significantly less effective than all other methods ($p < 0.001$). No other statistically significant differences existed among other methods.

Additional details split up by age assurance completers and non-completers are shown in Figure 9 in Section A.2. Non-completers of age verification tended to say that age assurance was less effective. However, even among age assurance completers, the highest rated effectiveness was only slightly higher than among all participants: 47.5% for Gov-ID + Liveness and 48.1% for Gov-ID with no reassurance (NR).

In open-ended responses, study participants commented on effectiveness. Some participants explained why they felt the checkbox was ineffective, for example, "There is zero actual verification [for the checkbox], its basically the honor system." Some participants expressed skepticism about ID systems' effectiveness. One participant wrote, "I guess it would be easy for some people to just use their parents ID, while other people would probably follow the rule." Another wrote, "It probably will only deter someone who is just slightly motivated to access the content, while increasing the risk of people getting used to providing their ID all the time." Some commenters voiced a belief that AI tools would make it easier to bypass even ID-based age assurance systems. One participant wrote, "Chat gpt provide me an over 18 CA drivers license." Some participants expressed that liveness checks would be difficult, but not impossible to bypass. One participant wrote, "A liveness check will be hard to spoof, but I'm sure workarounds will be found."

2.4 ID methods were seen as more inconvenient

Convenience was measured by asking the question, "How convenient was it for you to complete the age verification process?" to those who completed the age verification step, and collecting responses on a five-point Likert scale.

Results and significance are shown in Figure 8. Convenience ratings varied significantly by condition ($\chi^2 = 116.71$, $p < 0.001$), with checkbox self-attestation rated as more convenient than all other methods (pairwise comparisons: all $p \leq 0.03$). Among the government-ID conditions, no significant differences were found between the no-reassurance, simple-reassurance, and compound-reassurance conditions (all $p \geq 0.12$). The government-ID with liveness check condition was rated as significantly less convenient than both the AI and email-based age estimation methods (both $p < 0.001$). The AI and email-based conditions did not differ significantly from each other ($p = 1.00$), with both rated as more convenient than most government-ID approaches but less convenient than checkbox self-attestation.

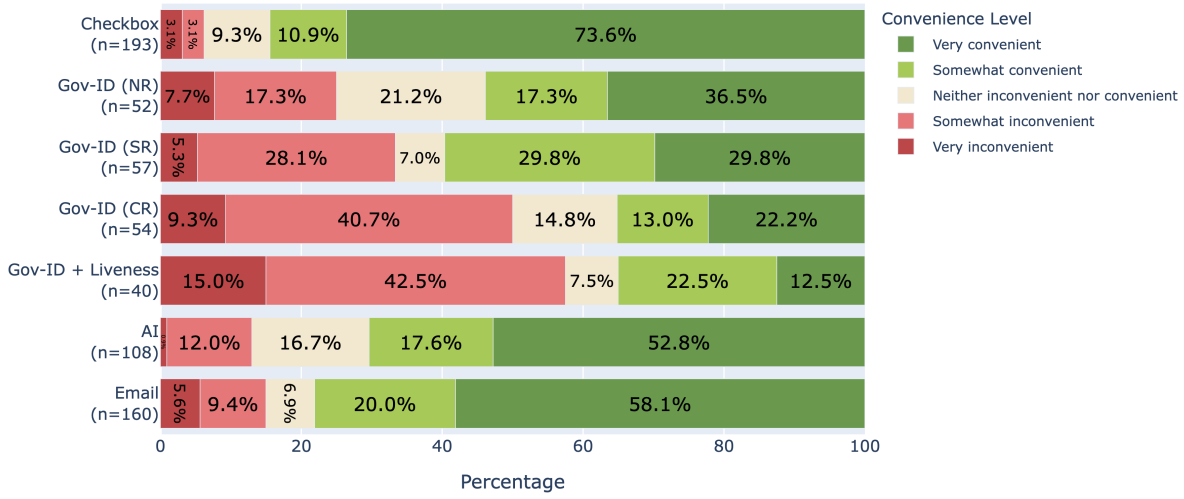

Figure 8. Perceived **convenience** among those who completed age verification. Pairwise comparisons of perceived convenience across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values) shown in Table 11 found statistically significant differences between the Checkbox condition and all Gov-ID conditions (including Gov-ID + Liveness) at $p < 0.001$, and between Checkbox and AI at $p < 0.01$. Gov-ID + Liveness was significantly less convenient than all other conditions except Gov-ID (CR) at $p \leq 0.05$.

Although we did not ask a free-response question about convenience, some participants commented on convenience in a more general free-text response. Some participants described intrusion or annoyance of some forms of age verification because of the time it takes, or needing to to take extra steps to verify. One participant wrote, "Age verification is already a barrier in that it takes a while to complete and the idea of scanning personal documents might put one off." One participant in the AI condition wrote, "to take the picture I'd have to get out of bed and turn on the lights which I didn't want to do (at 3am)" – a concern that would also apply in the ID settings. A participant in the Gov-ID condition wrote about the annoyance of dealing with the process itself and failures in the system: "It's also annoying to have to take out my id, take a picture, and potentially deal with an automated system not recognizing it." Another wrote that email verification seemed less inconvenient than other approaches: "Age verification is typically pretty intrusive. I use email to sign up for things all the time, so its not much of an inconvenience to use it for age verification." Some participants commented on AI systems specifically. One participant wrote, "I'm 19 and everyone tells me I look young. I was preparing myself to be pretty annoyed because I expected that the AI would say I was underage even though I'm not."

2.5 Additional findings

Confusion over who gets the data. In the survey, we asked people, "During the verification process, who did you think collected your age data? Select all that apply." As shown in Table 2, there was some confusion on this point. Most participants believed that an entity other than CMU was collecting their age data (in some cases, in addition to CMU). This shows that our deception was largely effective. However, as nearly half the participants thought that CMU would receive their age data as well, it suggests, that in real age verification situations, users may be concerned about content providers (not only age verifiers) receiving their personal information.

Table 2. Participants' perceptions of which entity collected their age data.

Response	Count (%)
AgeGuardian only	290 (32.92%)
CMU only	57 (6.47%)
AgeGuardian and CMU	340 (38.59%)
Not sure	112 (12.71%)

Percentages do not sum to 100% because some participants selected *"Other"* and described entities not listed in the table.

More comfortable with CMU. We also asked participants if their comfort was affected by knowing that the information was used by CMU. As shown in Table 3, 68.5% of participants felt at least somewhat more comfortable sharing this information with CMU than with others. This suggests that our completion rates and comfort responses are likely higher than they would be for an age verifier in the wild.

Table 3. Impact of knowing CMU is running the study on comfort with the age verification process.

Response	Count (%)
A lot more comfortable	332 (37.68%)
Somewhat more comfortable	272 (30.87%)
No change in comfort	230 (26.11%)
Somewhat more uncomfortable	25 (2.84%)
A lot more uncomfortable	22 (2.50%)

Question wording: "How did knowing the fact that CMU is running the research study site impact your comfort with the age verification process? It made me:"

Few participants read privacy policies. During the $n = 1635$ web experiment, fourteen study participants (0.86%) clicked on the privacy policy of AgeGuardian.

A Appendix

This Appendix contains summary statistics, as well as additional figures and tables.

A.1 Summary Demographics

Overall we collected $n = 1635$ valid responses to our web experiment and $n = 881$ valid responses to our follow-up survey after removing duplicates and suspected bots. See Tables 4, 5, and 6 for more detail. We had respondents from all 50 states as shown in Table 7.

Table 4. Gender of study participants who completed the survey

Gender	Count (%)
Man	609 (69.13%)
Woman	164 (18.61%)
Non-binary	45 (5.11%)
Self-describe	29 (3.29%)
Prefer not to disclose	34 (3.85%)
Total	881 (100%)

Table 5. Education levels of study participants who completed the survey

Education level	Count (%)
Bachelor's degree	312 (35.41%)
Graduate or professional degree	207 (23.50%)
Some college, but no degree	197 (22.36%)
Associates or technical degree	71 (8.06%)
High school diploma or GED	55 (6.24%)
Prefer not to respond	20 (2.27%)
Some high school or less	19 (2.16%)
Total	881 (100%)

Table 6. Race/Ethnicity study participants who completed the survey

Race/ethnicity	Count (%)
White	661 (75.03%)
Asian	73 (8.29%)
Black or African-American	14 (1.59%)
American Indian/Native American or Alaska Native	5 (0.57%)
Native Hawaiian or Other Pacific Islander	1 (0.11%)
Hispanic/Latinx	7 (0.79%)
Other	72 (8.17%)
Prefer not to say	48 (5.45%)
Total	881 (100%)

Table 7. States of residence of study participants who completed the survey

State	Count (%)
California	109 (12.37%)
Pennsylvania	63 (7.15%)
Texas	49 (5.56%)
Washington	48 (5.45%)
Illinois	43 (4.88%)
Michigan	37 (4.20%)
Ohio	36 (4.09%)
Massachusetts	34 (3.86%)
Virginia	34 (3.86%)
New York	33 (3.75%)
Utah	31 (3.52%)
Oregon	28 (3.18%)
Wisconsin	25 (2.84%)
Florida	24 (2.72%)
New Jersey	22 (2.50%)
Colorado	21 (2.38%)
Minnesota	18 (2.04%)
North Carolina	18 (2.04%)
Georgia	15 (1.70%)
Arizona	14 (1.59%)
Maryland	13 (1.48%)
Indiana	12 (1.36%)
Tennessee	12 (1.36%)
Missouri	12 (1.36%)
Idaho	8 (0.91%)
Kentucky	8 (0.91%)
Connecticut	8 (0.91%)
Montana	7 (0.79%)
Alabama	7 (0.79%)
Kansas	7 (0.79%)
Maine	6 (0.68%)
Washington, D.C. (Federal District)	6 (0.68%)
Iowa	6 (0.68%)
Nevada	6 (0.68%)
Arkansas	5 (0.57%)
South Carolina	4 (0.45%)
Oklahoma	4 (0.45%)
Nebraska	4 (0.45%)
Louisiana	3 (0.34%)
Hawaii	3 (0.34%)
Rhode Island	2 (0.23%)
New Mexico	2 (0.23%)
New Hampshire	2 (0.23%)
Mississippi	2 (0.23%)
West Virginia	2 (0.23%)
Vermont	2 (0.23%)
Delaware	2 (0.23%)
Alaska	2 (0.23%)
Wyoming	1 (0.11%)
North Dakota	1 (0.11%)
Prefer not to respond	20 (2.27%)
Total	881 (100%)

A.2 Additional Tables and Statistics

Statistics for completion rates. See Table 8 for additional statistics about completion rates.

Table 8. Pairwise comparisons of completion rates (Wilcoxon rank-sum tests); Holm-Bonferroni-adjusted p-values).

Comparison	Adjusted p
Checkbox vs. Gov-ID (NR)	$< .001$
Checkbox vs. Gov-ID (SR)	$< .001$
Checkbox vs. Gov-ID (CR)	$< .001$
Checkbox vs. Gov-ID + Liveness Check	$< .001$
Checkbox vs. AI	$< .001$
Checkbox vs. Email	$< .001$
Gov-ID (NR) vs. Gov-ID (SR)	1.00
Gov-ID (NR) vs. Gov-ID (CR)	1.00
Gov-ID (NR) vs. Gov-ID + Liveness Check	0.79
Gov-ID (NR) vs. AI	$< .001$
Gov-ID (NR) vs. Email	$< .001$
Gov-ID (SR) vs. Gov-ID (CR)	1.00
Gov-ID (SR) vs. Gov-ID + Liveness Check	0.59
Gov-ID (SR) vs. AI	$< .001$
Gov-ID (SR) vs. Email	$< .001$
Gov-ID (CR) vs. Gov-ID + Liveness Check	0.13
Gov-ID (CR) vs. AI	$< .001$
Gov-ID (CR) vs. Email	$< .001$
Gov-ID + Liveness Check vs. AI	$< .001$
Gov-ID + Liveness Check vs. Email	$< .001$
AI vs. Email	$< .001$

Note: NR = no reassurance; SR = simple reassurance (data used only for age verification); CR = compound reassurance (data used only for age verification and guarantee of deletion of identifying information after access)

Statistics for comfort. See Table 9 for additional statistics on self-reported comfort across conditions.

Table 9. Pairwise comparisons of self-reported comfort across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values).

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

$^*p < .05$ (Holm–Bonferroni corrected).
Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 252.21$, $p < .001$).

Statistics for effectiveness. See Table 10 for additional statistics on perceived effectiveness across conditions.

Table 10. Pairwise comparisons of perceived effectiveness across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values).

Comparison	Adj. 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

$^*p < .05$ (Holm–Bonferroni corrected).
Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 62.37$, $p < .001$).

Effectiveness results by completion. See Figure 9 for a graph of user-perceived effectiveness split by whether the participant completed age verification.

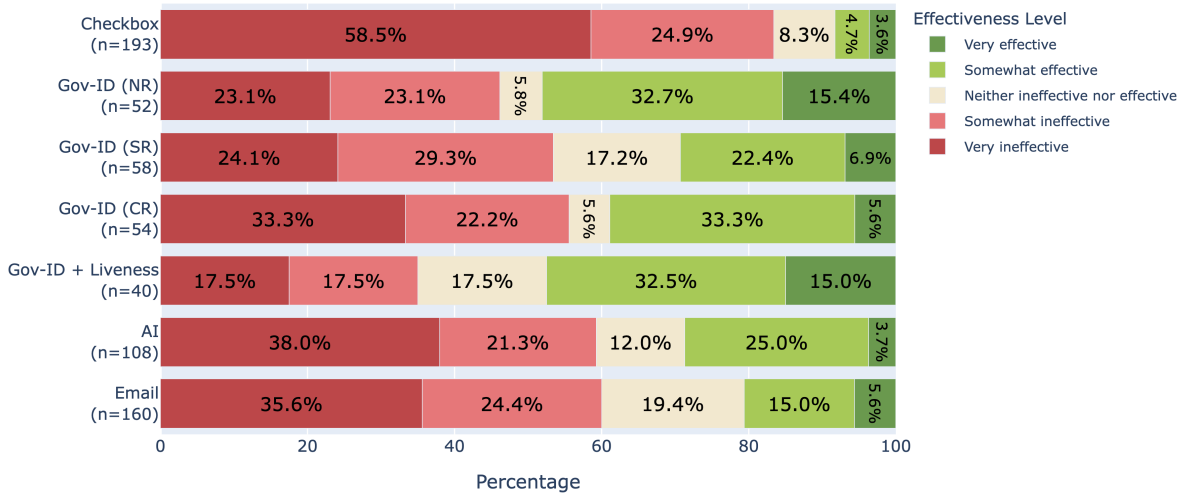

(a) Reported effectiveness among participants who **completed** age assurance

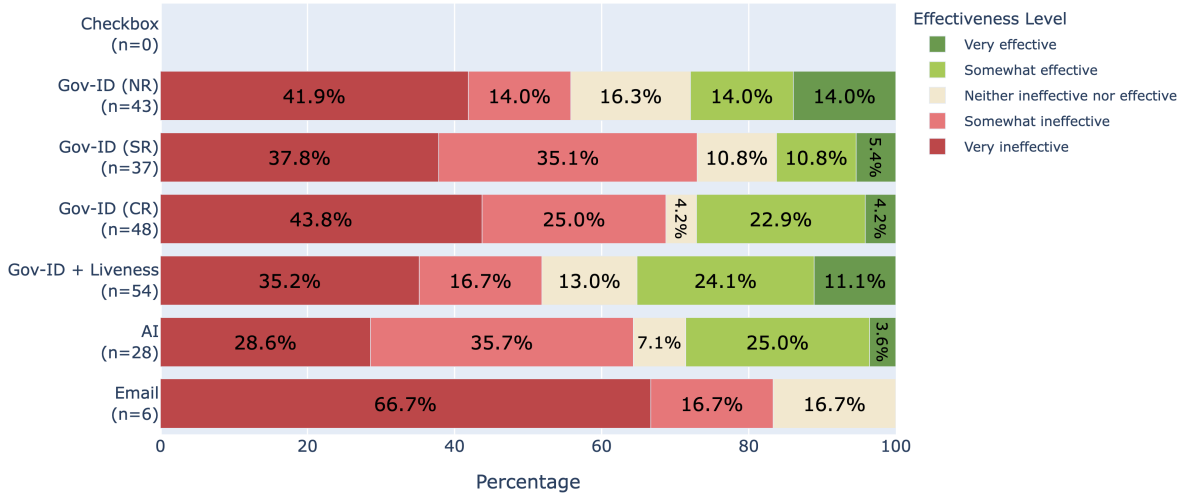

(b) Reported effectiveness among participants who **did not complete** age assurance

Figure 9. Reported effectiveness among participants who **completed** age assurance (top) and those who **did not complete** age assurance (bottom)

Statistics for convenience. See Table 11 for additional statistics on perceived convenience across conditions.

Table 11. Pairwise comparisons of perceived convenience across conditions (Wilcoxon rank-sum tests with Holm–Bonferroni-adjusted p-values).

Comparison	Adj. 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

$^*p < .05$ (Holm–Bonferroni corrected).
Overall differences were significant (Kruskal–Wallis $\chi^2(6) = 116.71$, $p < .001$).