Extended abstract: Measuring User Responses to Online Age Verification

Yanzi Veronica Lin, Cheng Zhang, Madelyne Xiao, Weiqian Zhang, Vivianna Lieu, Wenchao Hu, Lorrie Faith Cranor, and Sarah Scheffler
Carnegie Mellon University

December 2025

The U.S. Supreme Court's 2025 decision in Free Speech Coalition v. Paxton established that age verification systems must be "adequately tailored" to prevent undue burdens on adults' First Amendment rights. This study empirically examines how different age assurance methods (e.g. government-issued ID, AI facial estimation, checkbox self-attestation) and accompanying data handling disclosures influence user behavior, yielding an empirical lens through which one can measure the burdens and preferences of users with regard to age verification methods. Specifically, we conduct a deceptive web experiment ($n = 1264$ at time of writing) to measure users' completion rates for a randomly-chosen variant of age verification; further, we follow up with participants with a survey ($n = 651$ at time of writing). These user behaviors should be taken into account when determine what constitutes "adequate tailoring" of age verification systems. Moreover, as many U.S. states are transitioning toward digital ID, these results should be considered for a smooth roll-out of future systems for verifying things other than age as well.

In our study design, we deceptively informed participants that the purpose of the study was to research romantic media. After obtaining consent, the user would be informed that we must verify their age before conducting the study. They would be directed to a mock age verification website called AgeGuardian, which was also a website set up by the researchers. AgeGuardian would run the user through a mock (but seemingly real) randomly-selected form of age verification: checkbox self-attestation, taking a picture of a government-issued ID, using AI facial age estimation, or an email lookup. Additional variants tested different "reassurance statements" accompanying the government-ID conditions. (See Figure 1 for some mockup pictures of AgeGuardian.) The AgeGuardian website was programmed to leave all sensitive information on the user's device; even though the website would appear to upload their ID or face picture, it was never actually sent over the network and was never seen by the researchers. If the user age verified, they would be directed to a debrief and survey. If the user did not complete age verification (clicked "Exit Study" button, closed tab, or timed out) then the user would be sent an email to the address they used for the consent form, which would provide them with a debrief, inform them that their participation in Part 1 was complete, give them options to remove their data, and provide them with a link to the Part 2 survey to complete if they wished. Our participants are recruited using ads on Reddit, Meta, and TrafficJunky (an advertising site putting ads on age-verified pornography websites).

This study follows up from an earlier preliminary study, which we presented at the IAB/W3C Workshop on Age Verification.[1] That preliminary study had fewer participants ($n = 99$), investigated fewer age verification methods, and used a different deceptive setup (a gambling website).

Results for the web experiment portion of the study are shown in Table 1 To summarize our results, the checkbox had the highest completion report and user-reported comfort. Methods involving government-issued ID verification had lower completion rates and lower comfort. Data handling disclosures generally improved completion rates, but a small amount compared to overall results. Our team is still reviewing qualitative results, but results indicate (1) both support and non-support for age verification in general, (2) concerns about privacy and surveillance, and (3) concerns about security, with participants expressing reluctance to share sensitive information with unfamiliar entities.

[1] IAB/W3C Workshop on Age Verification: `https://datatracker.ietf.org/group/agews/about/`, see our slides and paper which presented a preliminary version of this study.

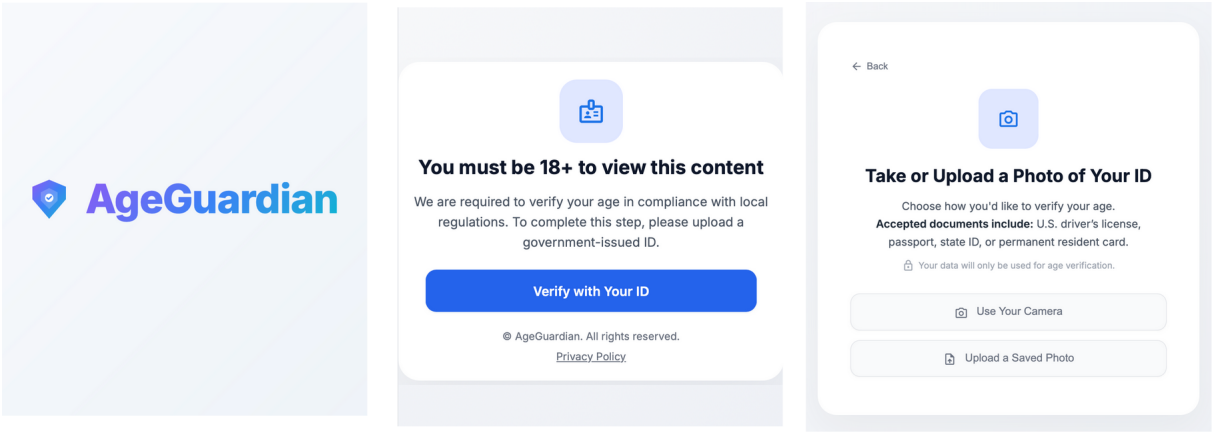

Figure 1: Screenshots of the researcher-controlled mock age verification website, AgeGuardian. Users were deceptively instructed to age-verify with AgeGuardian before conducting a "romantic media" study. After completing (or not completing) age verification, the user would be directed to (or sent) a debrief informing them that the main purpose of the study was to measure whether or not they would (appear to) verify their age using their assigned age verification condition on AgeGuardian.

Condition	**Web Experiment Participants**	**Completed Age Verification** (%)
Checkbox self-attestation	168	166 (99%)
Gov-ID (no reassurance)	189	40 (21%)
Gov-ID (simple reassurance)	189	44 (23%)
Gov-ID (compound reassurance)	182	49 (27%)
Video selfie + Gov-ID	190	32 (17%)
AI Facial Age Estimation	175	90 (51%)
Email lookup	171	145 (85%)
Total	1264	

Table 1: Results from the web experiment portion of the study ($n = 1264$).