Knight
Georgetown
Institute

JULY 2026

Age Assurance Rules:

An Implementation Guide

Amy Winecoff
Alissa Cooper
Knight-Georgetown Institute

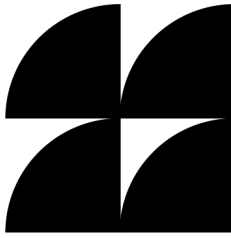

About the Knight-Georgetown Institute

The Knight-Georgetown Institute (KGI) is dedicated to connecting independent research with technology policy and design. KGI serves as a central hub for the growing network of scholarship that seeks to shape how technology is used to produce, disseminate, and access information. KGI is designed to provide practical resources that policymakers, journalists, and private and public sector leaders can use to tackle information and technology issues in real time. Georgetown University and the Knight Foundation came together to launch the institute in 2024. Learn more about KGI at https://kgi.georgetown.edu.

Table of Contents

I. Introduction

Age assurance, the process of determining whether a user is eligible to access age-restricted services, products, or features, is becoming a foundational component of policies intended to make online experiences safer for youth. While it can potentially play a role in a comprehensive approach to youth online safety, age assurance is not in and of itself a safety intervention and cannot address online harms on its own. Nor is it without costs. Age assurance systems introduce tradeoffs affecting privacy, users' ability to access online services, expression, competition, and the openness of the device ecosystem for users of all ages. Effective age assurance policy must therefore ensure that potential benefits are appropriately balanced against potential risks to users and the broader digital ecosystem.

The right balance will typically depend on context, including the type of service or product involved (e.g., social media, AI chatbots, adult websites), the harms the policy is intended to address (e.g., predation, addictive use, emotional manipulation), and the age-dependent restrictions intended to be imposed. The use cases for age assurance largely fall into two categories: (1) establishing safer defaults that provide tailored age-appropriate experiences for youth, and (2) blocking access to content, services, or accounts for users under a specific age.

Across the US and around the world, age assurance mandates supporting both use cases have been proposed and adopted in a variety of forms. A common sequence in many jurisdictions has seen the policymaking process start with high-level requirements being incorporated into new laws or legislative directives, with more detailed regulations, implementation requirements, or rules being defined afterward (if they get defined at all). In some cases the enacting laws themselves may include significant implementation detail, but many legislative proposals recognize that the complexity of age assurance necessitates having an expert regulator, state attorney general's office, or public authority construct more detailed rules than what is typical in legislation.

This implementation guide is intended for policy stakeholders grappling with how to design those detailed rules after a more general legal mandate that requires age-based blocking or defaults has been put in place. The recommendations below may also be useful in cases where implementation details are being proposed in legislation, or as part of settlements or court-ordered remedies in litigation. Building from *Age Assurance Online: A Technical Assessment of Current Systems and Their Limitations*, the Knight-Georgetown Institute's technical assessment of existing age assurance technologies, the recommendations focus on how rules can support age assurance regimes that create safer online experiences for youth, while acknowledging the serious tradeoffs presented by the mandatory use of these systems.

The next section defines a few key terms and concepts. The recommendations that follow are organized around the major components of a balanced age assurance regime—from foundational rule design and operational requirements to evaluation, reporting, and privacy protections.

Age assurance may constitute one component of a broader, comprehensive approach to youth safety online, but no single system can simultaneously optimize for accuracy, privacy, service availability, and circumvention resistance. Every age assurance system will also unavoidably introduce some degree of friction for adults and eligible minor users of online services. Rules should therefore be designed to maximize safety benefits to youth while minimizing burdens on other users and the broader digital ecosystem. The recommendations that follow are designed to help policymakers strike this balance.

II. Terminology

Covered providers. This implementation guide assumes that age assurance requirements are being imposed on **covered providers** that operate online applications or services. These may include app developers, website operators, service providers or platforms that offer specific products or services (e.g., social media, AI chatbots, etc.).

Properties of age assurance systems. There are four key criteria involved in the technical evaluation of any age assurance system:

- **Privacy:** the degree to which user information is revealed that otherwise would not be in the absence of age assurance;
- **Availability:** ability for the eligible population (e.g., adults or minors who are age-eligible) to gain access;
- **Accuracy:** correctness of age determinations in the absence of circumvention attempts; and
- **Circumvention resistance:** robustness against attempts by users to establish an age different from their true age.

Age signals. Age assurance methods based on different **age signals** have different properties when it comes to privacy, availability, accuracy, and circumvention resistance. Common age signals used to determine user ages include:

- Self-declaration, or parental declaration of a minor's age
- Commercial and government records, consisting of banking records, mobile network operator records, credit cards, or other commercial and government records retrieved by name or email address
- Physical government ID or digital IDs
- Facial age estimation
- Behavioral signals, consisting of data that covered providers otherwise already collect in the course of offering a product or service

Accuracy. When evaluating accuracy of age assurance systems, two values are important:

- The **false rejection**[1] rate, representing the fraction of users within the eligible age range (e.g., over 18) who are not permitted to access the system, for instance.
- The **false acceptance** rate, representing the fraction of users outside the eligible age range (e.g., under 18) who are permitted to access the system.

All of these terms are used throughout the rest of the guide.

III. Rule Design

This section offers three recommendations for how to design a set of rules that, collectively, integrates and balances youth safety goals with other goals and properties of age assurance systems.

1. **Policy objective: Complete circumvention resistance should not be a policy objective.**

No age assurance system can prevent all underage users from accessing restricted services or experiences without imposing substantial costs on everyone. Achieving that level of restriction requires collecting more sensitive and personally identifiable data, demanding repeated assessment that compounds friction and privacy risks, or deploying highly restrictive device-level controls that will impede legitimate users' access. Complete circumvention resistance may also deny parents the ability to authorize their children to access age-gated services, products, or features at their discretion.

A balanced policy goal is to meaningfully reduce risk through a multifaceted youth online safety approach, not to achieve perfect enforcement through age assurance. What it means to meaningfully reduce risk to youth depends heavily on context, including both the harms that an age assurance mandate is intended to mitigate (e.g., sleep interruptions from nighttime notifications or prevalence of suicidal ideation) as well as the technology being targeted (e.g., AI companions versus short-form video platforms). The laws and policies that lead to the implementation of age assurance rules should serve measurable harm-reduction goals, rather than making perfect age-gating the policy goal.

2. **Responsibility for covered providers: Rules (and the laws that call for them) should assign responsibility for complying with youth safety requirements to covered providers, rather than third parties that provide age signals or age assurance.**

Covered providers control the features and design decisions that youth online safety policies seek to influence, and they are therefore best positioned to implement age-appropriate protections. They should be accountable for meeting legal requirements for youth online safety, including the adoption

[1] The statistics and testing literature uses a number of terms for error rates, including sensitivity versus specificity, false positive versus false negative, and type I versus type II errors. These can often be hard to interpret because of confusion about whether a “positive” result leads to acceptance or rejection. This guide uses the terms “false rejection” and “false acceptance” for clarity.

or reliance on age assurance systems sufficient to meet those requirements, regardless of whether they perform age assurance themselves or rely on operating systems, third-party age assurance providers, or other parties to assess age or transmit age signals.

This approach preserves flexibility for covered providers to adopt different technical approaches to age assurance while maintaining clear responsibility for youth safety outcomes. Operating systems or other entities that may be implicated in the age assurance process can potentially support age assurance in limited ways, as described in KGI's companion piece, *Assigning Responsibility for Age Assurance: Recommendations for Youth Online Safety*. But covered providers that control how users access and interact with their products and services should have ultimate responsibility for youth safety.

To reduce the burden on small providers and ease enforcement, legal obligations for age assurance could be confined to services that reach a large number of users or pose the greatest risks to youth. This narrower scope both cuts unnecessary collection of age-related information, particularly by websites or app developers without extensive resources for implementing the highest level of privacy protections, and focuses regulatory efforts where they are likely to have the greatest impact.

3. **Safety requirements tied to age assurance: Age assurance obligations should be imposed together with safety requirements that obligate covered providers to act on the basis of age determinations.**

Age assurance is not itself a safety intervention. Rather, it is a mechanism that can be used to determine which age-based requirements apply (e.g., whether to apply strong privacy settings, to block nighttime notifications, or to disable long-term memory for AI chatbots). Legally mandating that providers transmit user ages or adopt age assurance without specifying the safety measures that covered providers must implement for youth does not make those providers' products and services safer. Age assurance requirements should therefore be imposed together with corresponding safety requirements designed to mitigate the harms the rules are intended to address.

IV. Operational Rules

This section provides five recommendations for how to implement technical and operational requirements related to selection of age assurance methods, accuracy, recourse for users who get denied access, and parental consent.

4. **Multiple methods: Covered providers should be required to offer multiple distinct age assurance methods to ensure services remain available to eligible users.**

No single age assurance method is simultaneously highly accurate, privacy protective, and available to all users. Some users lack government-issued identification, some cannot successfully complete biometric checks, and others may have legitimate privacy concerns about particular forms of age assurance. Requiring covered providers to offer multiple methods that rely on distinct age signals (self-declaration or parental declaration, commercial or government records, government IDs, facial age estimation, or behavioral signals) reduces the likelihood that access depends on the limitations of any single approach. Having more than one option available should allow users to choose the method that best aligns with their individual circumstances (e.g., having a specific disability, not owning a driver's license) and privacy preferences.

Offering multiple age assurance methods may increase opportunities for some users to attempt to circumvent restrictions. Nevertheless, because the objective of age assurance should be to reduce harm to youth rather than prevent all underage access, the benefits for allowing all users an opportunity to demonstrate their age can be balanced against this risk.

5. **Accuracy: Covered providers should be required to meet specified accuracy rates for each age assurance method they offer to users, as well as collective accuracy rates for their overall system.**

Some age assurance systems, especially age estimation systems, inherently have some level of error. Requiring covered providers to choose age assurance methods that meet specified accuracy thresholds (which can be determined by the rulemaking authority, preferably with expert consultation) will discourage covered providers from choosing methods that are known to be error-prone.

However, the accuracy of age assurance systems that route users across multiple methods cannot be inferred from evaluating individual methods alone. Understanding overall system accuracy requires covered providers to report end-to-end evaluation of real-world false acceptance and false rejection rates collectively across all the provider's supported methods. Rules should therefore include overall system accuracy requirements in addition to individual method accuracy thresholds.

6. **False rejections: Accuracy thresholds and enforcement penalties should apply for both false acceptances and false rejections.**

Many existing policy proposals focus primarily on preventing underage access while paying comparatively less attention to the consequences of incorrectly denying access for eligible users. Yet false rejections can exclude eligible users from legitimate access to apps and services. Accuracy rules should therefore account for both types of error, weighing harm-reduction benefits for young people against access impacts for eligible users. This ensures that providers are incentivized to prevent underage access without relying on overly restrictive systems that exclude eligible users.

7. **Recourse for false rejections: Covered service providers should be required to offer users denied access with timely and meaningful avenues for recourse.**

All age assurance systems will produce false rejections, meaning some eligible users will inevitably be turned away. Without effective recourse mechanisms, those users may lose access to important services, information, or opportunities for expression despite meeting all legal requirements.

Reasonable, time-bounded appeals processes should be a mandatory requirement for effective age assurance. Recourse mechanisms should adhere to the same privacy and availability objectives as the underlying age assurance system. Users should not be required to disclose substantially more personal data simply because an initial age assurance attempt produced a false rejection.

8. **Parental consent: Where parental consent is a required component of an age assurance mandate, it should be supported either through (1) device- or OS-level configurations that transmit signals to covered providers, or (2) covered providers establishing parental relationships via persistently linked accounts. Parental identity verification should not be required to confirm parental consent.**

Creating a highly accurate, circumvention-resistant parental verification system for use by services broadly available on the internet would require collecting identity information from both adults and minors, as well as additional information to prove parental relationships. Mandating parental identity verification therefore risks imposing substantial privacy, availability, and implementation burdens that are not warranted by the likely harm-reduction benefits, particularly when less invasive alternatives exist, even if those alternatives are not fully circumvention-proof. Requirements to support parental consent mechanisms should instead focus on facilitating simple configuration at the device, operating system,[2] or covered provider level rather than identity- and relationship-based verification.

V. Evaluation and Reporting Rules

[2] For a more detailed discussion of policy approaches to operating system-level signaling in age assurance beyond parental consent, see *Assigning Responsibility for Age Assurance: Recommendations for Youth Online Safety*.

Covered providers' self-reported metrics on how their systems perform for accuracy, privacy, or circumvention resistance should not be taken at face value. Policymakers and the public need reliable mechanisms for assessing whether age assurance implementations are accurate, effective, and compliant with applicable requirements. Because measurement of the performance of these systems depends both on the system design as well as the evaluation methodology, effective oversight requires clear requirements for measuring performance, transparency regarding evaluation methods and results, and independent verification of provider claims. This section provides recommendations about each of these three requirements. Together, these mechanisms can provide policymakers and the public with a more reliable basis for understanding how age assurance systems perform in practice.

9. **Evaluation transparency: Covered providers should be required to publicly report their evaluation methods and system performance results in sufficient detail to enable independent replication.**

Evaluation results are only as reliable as the methods used to produce them, and weak methodologies can artificially inflate performance estimates. Requiring covered providers to publicly disclose their evaluation methods – including either detailed descriptions of their evaluation data or a representative sample of the evaluation data – and results enables independent scrutiny and replication of performance claims.[3] This transparency incentivizes covered providers to use rigorous evaluation practices, to accurately report results, and to adhere to industry best practices.

For some performance qualities such as circumvention resistance, quantitative measures are often less useful, and the specific circumvention strategies that require testing are likely to evolve.[4] In these cases, covered providers should disclose qualitative descriptions of their circumvention resistance testing that is detailed enough to enable conceptual replication without creating additional security risks. Likewise, covered providers should describe their methods for testing availability since this quality does not necessarily lend itself to quantification.

Where covered providers rely on third parties to provide age signaling or age assurance services, covered providers should be responsible for reporting system performance using data that are representative of how they have deployed the third-party services. Covered providers that publish

[3] Replication can take multiple forms. At one end of the spectrum, assessors may reproduce results using the exact code and data employed by the service provider ("exact replication"). At the other, evaluators may independently implement the reported methods using data and procedures consistent with the provider's description ("conceptual replication"). In this context, conceptual replication is generally preferable. Requiring providers to disclose proprietary datasets may create unnecessary risks to user privacy. Moreover, conceptual replication provides a more meaningful test of whether reported performance is robust and reproducible. If an independent assessor cannot reproduce a provider's reported results within an acceptable margin of error, it may indicate flaws in the provider's methodology, inaccuracies or omissions in its reporting, or that the reported performance does not generalize beyond the provider's original evaluation setup. Regardless of the underlying cause, such discrepancies should trigger additional investigation and, where they cannot be resolved, appropriate enforcement action.

[4] For a more specific treatment of this issue in the context of a live rulemaking, see Rescorla and Cooper, "First Steps Toward Operationalizing Age Assurance Mandates: New York SAFE for Kids Act Proposed Rules."

third-party performance assessments to meet their own evaluation transparency obligations should remain responsible for ensuring the completeness and accuracy of those assessments.

10. **<u>Performance by subpopulation:</u> Covered service providers should be required to evaluate and report system performance separately for relevant demographic groups**.

Age assurance systems are unlikely to perform equally well for all users. Differences in false acceptance rates, false rejection rates, or other performance measures between demographic groups[5] may result in some populations being disproportionately denied access or subjected to greater risks where fallback mechanisms require more personal data or more privacy-invasive forms of age assurance. Mandating that covered service providers assess and publicly report evaluation results per relevant demographic group helps external stakeholders understand the broader societal impacts of age assurance systems that may not be apparent from overall performance metrics. Even where regulations do not establish nondiscrimination requirements, public reporting of population subgroup performance metrics can encourage covered providers to identify, investigate, and reduce disparities in performance between demographic groups.

11. **<u>Independent verification</u>: Covered providers should be required to undergo assessment by an independent, qualified third party to verify reported performance claims.**

Voluntary transparency from covered providers about their (or their vendors') results is helpful but not sufficient. Independent verification is necessary to verify that company-reported results are accurate and that systems meet all applicable requirements. For that independence to be meaningful, assessors must be genuinely qualified to assess these systems and held to external professional standards that ensure conflicts of interest are appropriately managed.[6]

Granting assessors sufficient access is equally important: age assurance systems are often complex, relying on external services or layered technical infrastructure. Without access to all relevant data and parts of the stack, independent assessors may be unable to fully assess whether the system as a whole meets applicable requirements.

[5] Performance of facial estimation methods for age assurance have been shown to vary as a function of both algorithm and demographic characteristics. See Grother et al., *Face Analysis Technology Evaluation*.
[6] Enabling an ecosystem of independent evaluators requires a range of supporting policies and institutional conditions. For a discussion of these factors in the context of AI assurance, see Groves et al., *Going pro? Considerations for the Emerging Field of AI Assurance*.

VI. Privacy Rules

Age assurance systems collect, transmit, and process information about users, creating privacy risks that should be minimized wherever possible. The five recommendations for privacy rules below focus on limiting the disclosure of age-related information, minimizing unnecessary collection and reassessment, and encouraging technical approaches that reduce data exposure by design.

12. **Disclosure limitation: Age signals should only be shared with services that have a legal obligation to act on them, or with entities acting on those services' behalf. They should be transmitted only when requested by covered providers and only to the extent necessary to assess and communicate age eligibility or age ranges.**

Age-related information, even when shared as age ranges, can be revealing when combined with other sources of data. Limiting disclosure to services that actually need the information, reducing the precision of what is shared, and constraining transmission to on-demand requests all help reduce opportunities for profiling, tracking, and secondary uses. Privacy rules should therefore require that no information beyond what is necessary to communicate a user's age eligibility be disclosed to covered providers and the third-party service providers on which they rely.

Consistent with long-standing data security best practices, rules should also require that any data collected for age assurance be encrypted and secured using industry-standard practices to help protect against unauthorized access or interception.

13. **Data minimization: Covered service providers should collect only the information necessary to determine age eligibility and retain it only long enough to support a recourse process.**

Collecting, retaining, or combining data beyond what is necessary to determine age eligibility increases privacy and security risks while providing no corresponding benefit. Rules should require covered service providers to delete any data collected specifically for age assurance that is not necessary to maintain the resulting age-eligibility status or support a time-limited recourse process. For age assurance systems that rely on behavioral signals – data the covered provider otherwise collects anyway to provide its service or product – the data used for age determinations may be retained while the covered provider is using it to deliver the service or product.

Rules should prohibit information collected during the age assurance process from being combined with other personal data about the user except where necessary to determine age eligibility. Limiting collection, retention, and data linkage reduces opportunities for profiling and secondary use.

14. **<u>Purpose limitation</u>: Data collected specifically for age assurance should not be retained, repurposed, or shared with other entities for any purpose other than to complete the age assurance process.**

Allowing age assurance data to be retained or used for purposes beyond determining age eligibility creates incentives for unnecessary data collection and retention, increasing privacy risks and undermining user trust. Privacy rules should therefore strictly limit the collection, retention, and use of personal data to what is necessary to conduct the age assurance process and determine age-based eligibility. For systems that rely on behavioral signals, these limits can mirror limits on usage of the data to deliver the product or service, since the behavioral signals serve both age assurance and product purposes. Personal data collected specifically for age assurance should not be shared with entities other than those directly involved in conducting the age assurance process.

15. **<u>Technical privacy protections:</u> Rules should require the use of technical privacy mechanisms, including zero-knowledge proofs and selective disclosure, where applicable**.

Privacy protections are stronger when they are guaranteed through technical mechanisms rather than governance commitments alone. Technical approaches such as zero-knowledge proofs and selective disclosure enable eligibility determinations while cryptographically reducing the personal data disclosed. Zero-knowledge proofs allow a user to prove that they meet a requirement, such as being over 18, without revealing other personal data. Selective disclosure allows a user to share only the specific information needed for a particular purpose, rather than an entire identity document or account record. The circumstances under which each of these mechanisms are suitable depend on the architecture of the age assurance system in use.

As these approaches mature and become more commercially available,[7] implementation rules should require that covered providers adopt methods that incorporate these techniques, where the system architectures involved are suitable for these techniques to be deployed.

VII. Conclusion

Age assurance is a tool, not a solution. Used thoughtfully and governed well, it may help reduce risks to young people online by underpinning specific safety mitigations. Used poorly or designed without regard for the burdens it imposes, it can harm the very people it is intended to protect and the broader digital ecosystem. By drawing on the recommendations in this guide, policymakers can better ensure that age assurance rules reflect the realities of current technologies and appropriately account for both the harms they seek to mitigate and the tradeoffs they introduce.

[7] Such approaches are already beginning to be deployed (See e.g., *Linarducci, "Bringing Secure Digital Identity and Payment Tools to More People"; European Commission, "The EU Approach to Age Verification.")*

As the technology evolves and the evidence base on the impacts of age assurance develops,[8] policymakers should remain focused on whether age assurance systems are delivering meaningful benefits to youth online safety while preserving other important values, including privacy, service availability, competition, and openness. Rules established today may need to evolve and change as our collective understanding of age assurance and its impacts develops over time.

[8] Findings from academic researchers, independent regulators, civil society organizations, and government-commissioned evaluations have begun to illuminate the impacts of age assurance. Existing work has examined effects on platform dynamics (Agarwal et al., "Are You 18?"); the efficacy, robustness, privacy, and security of age assurance systems (Age Check Certification Scheme, *Age Assurance Technology Trial— Final Report*; Minocha et al., "Papers, Please"; Internet Matters and BMG Research, *The Online Safety Act*; Figueira et al., "Actions Speak Louder Than Chats."); proposed frameworks for designing and governing age assurance (Liu and Scheffler, "Adequately Tailoring Age Verification Regulations"; Collins, *Getting Age Assurance Right: A Risk-Based Framework for High-Risk Online Features*; Lawrence et al., "Tort Liability for Failure to Age Gate."); impacts on users and families (Bursztyn et al., "Why Bans Fail"; Lin et al., "Measuring User Responses to Online Age Verification Mechanisms Through Deceptive Experiment"; Barnes et al., "Assessing Early Effects of Australia's Social Media Minimum Age Act on Adolescents' Social Media Use.") and the overall effectiveness of age assurance policies and their implementation (Ofcom, *Use of Age Assurance Report 2026*).

Bibliography

Agarwal, Saharsh, Uttara M. Ananthakrishnan, Leonardo Madio, Matthew F. Mitchell, Martin Quinn, and Carlo Reggiani. "Are You 18? Age Verification and Adult-Only Consumption in the US." SSRN Scholarly Paper No. 6619018. Social Science Research Network, April 21, 2026. https://doi.org/10.2139/ssrn.6619018.

Age Check Certification Scheme. *Age Assurance Technology Trial— Final Report*. Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, Australian Government, 2025. https://www.infrastructure.gov.au/department/media/publications/age-assurance-technology-trial-final-report.

Barnes, Courtney, Alix Hall, Stephanie Mantach, et al. "Assessing Early Effects of Australia's Social Media Minimum Age Act on Adolescents' Social Media Use: Observational Study." Research. *BMJ* 393 (June 2026): e363695. https://doi.org/10.1136/bmj-2026-363695.

Bursztyn, Leonardo, Angela L. Duckworth, Rafael Jiménez-Durán, et al. "Why Bans Fail: Tipping Points and Australia's Social Media Ban." Working Paper No. 35162. Working Paper Series. National Bureau of Economic Research, April 2026. https://doi.org/10.3386/w35162.

Collins, Sara. *Getting Age Assurance Right: A Risk-Based Framework for High-Risk Online Features*. Public Knowledge, 2026. https://publicknowledge.org/policy/getting-age-assurance-right-a-risk-based-framework-for-high-risk-online-features/.

European Commission. "The EU Approach to Age Verification." May 13, 2026. https://digital-strategy.ec.europa.eu/en/policies/eu-age-verification.

Figueira, Olivia, Pranathi Chamarthi, Tu Le, and Athina Markopoulou. "Actions Speak Louder Than Chats: Investigating AI Chatbot Age Gating." arXiv:2602.10251. Preprint, arXiv, February 10, 2026. https://doi.org/10.48550/arXiv.2602.10251.

Grother, Patrick J., Mei Lee Ngan, Joyce Yang, George W. Quinn, and Austin Hom. *Face Analysis Technology Evaluation: Age Estimation and Verification*. NIST IR 8525. National Institute of Standards and Technology (U.S.), 2024. https://doi.org/10.6028/NIST.IR.8525.

Groves, Lara, Amy Winecoff, and Miranda Bogen. *Going Pro? Considerations for the Emerging Field of AI Assurance*. Ada Lovelace Institute, 2025. https://www.adalovelaceinstitute.org/report/going-pro/.

Internet Matters, and BMG Research. *The Online Safety Act: Are Children Safer Online?* Internet Matters, 2026. https://www.internetmatters.org/hub/research/online-safety-act-report-2026/.

Lawrence, Matthew B., Brett Frischmann, and Avi Sholkoff. "Tort Liability for Failure to Age Gate: A Promising Regulatory Response to Digital Public Health Hazards." *Journal of Tort Law* 18, no. 1 (2025): 283–309. https://doi.org/10.1515/jtl-2025-0018.

Lin, Yanzi, Cheng Zhang, Madelyne Xiao, Lorrie Faith Cranor, and Sarah Scheffler. “Measuring User Responses to Online Age Verification Mechanisms Through A Deceptive Experiment.” Paper presented at 10th Workshop on Technology and Consumer Protection (ConPro ’26). May 21, 2026. https://www.ieee-security.org/TC/SPW2026/ConPro/papers/lin-conpro26.pdf.

Linarducci, P. J. “Bringing Secure Digital Identity and Payment Tools to More People.” *Google*, June 4, 2026. https://blog.google/products-and-platforms/platforms/google-pay/secure-identity-payment-tools/.

Liu, Shuang, and Sarah Scheffler. “Adequately Tailoring Age Verification Regulations.” Paper presented at CSLAW ’26: Proceedings of the Symposium on Computer Science and Law. Version 2. March 3, 2026. https://dl.acm.org/doi/abs/10.1145/3788646.3789526.

Minocha, Shreyas, Isaac Sheridan, Harry Oppenheimer, Paul Pearce, and Michael A. Specter. “Papers, Please: A First Look at Age Verification on the Web.” *2026 IEEE Symposium on Security and Privacy (SP)*, May 2026, 4185–202. https://doi.org/10.1109/SP63933.2026.00161.

Ofcom. *Use of Age Assurance Report 2026*. 2026. https://www.ofcom.org.uk/online-safety/protecting-children/use-of-age-assurance-report-2026.

Rescorla, Eric, and Alissa Cooper. “First Steps Toward Operationalizing Age Assurance Mandates: New York SAFE for Kids Act Proposed Rules.” *Knight-Georgetown Institute*, March 12, 2026. https://kgi.georgetown.edu/research-and-commentary/first-steps-toward-operationalizing-age-assurance-mandates-new-york-safe-for-kids-act-proposed-rules/.

Winecoff, Amy, and Alissa Cooper. *Assigning Responsibility for Age Assurance: Recommendations for Youth Online Safety*. Knight-Georgetown Institute, 2026. https://kgi.georgetown.edu/research-and-commentary/assigning-responsibility-for-age-assurance