Knight Georgetown Institute

DECEMBER 1, 2025

New York Office of the Attorney General's Stop Addictive Feeds Exploitation (SAFE) for Kids Act Notice of Proposed Rulemaking

Knight-Georgetown Institute (KGI) Comments

Eric Rescorla
Alissa Cooper
Knight-Georgetown Institute

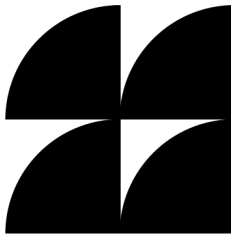

About the Knight-Georgetown Institute

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

The Knight-Georgetown Institute (KGI) is dedicated to connecting independent research with technology policy and design. Based at Georgetown University in Washington, D.C., 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.

We welcome the opportunity to provide input on the Notice of Proposed Rulemaking (NPRM) issued by the Office of the Attorney General (OAG) regarding the Stop Addictive Feeds Exploitation (SAFE) for Kids Act.[1] While we have done extensive work on policy approaches to improving user experiences with algorithmic feeds online,[2] that work is more relevant to the text of the law rather than the proposed rules. As such, our comments on the proposed rules focus on the age assurance portions, and in particular on operational and practical aspects.

An increasing number of jurisdictions across the country and the world are evaluating and adopting age assurance requirements of different kinds. This represents a sea change from how online services have been accessed over many decades, and it implicates a variety of important concerns and values for consumers, both adults and youth. The OAG's proposed rules represent the most significant attempt by a US public authority to specify how age assurance must be operationalized and monitored. This makes the OAG's proposed rules particularly salient for reflecting on the trade-offs involved in mandatory age assurance, including those between accuracy, privacy, circumvention resistance, service availability across the user population, and the production and consumption of speech and information. All age assurance systems must confront these trade-offs.

Given the nascency of mandatory age assurance regimes across jurisdictions, there has been little empirical observation or documentation of the efficacy and impact of the deployment of large-scale mandatory age assurance. It is critical for independent experts to be able to evaluate the social, technical, business, and individual impacts of mandatory age assurance. Every new age assurance regime presents an opportunity to build in structures and mechanisms that will allow for rigorous evaluation, research, and assessment to help inform the future shape of policy development. The OAG's proposed rules are the most meaningful opportunity to accomplish this in the US to date.

Where age assurance is being mandated, its design and implementation should reflect the currently understood best practices for protecting privacy and ensuring service availability. To that end, KGI's comments recommend changes to the proposed rules to make them technically sound and to better account for the properties of age assurance systems. We suggest ways to improve how the rules balance accuracy with availability given high error rates for some age assurance methods. We

[1] New York Office of the Attorney General, "Notice of Proposed Rulemaking."
[2] Moehring et al., "Better Feeds."

recommend moving circumvention detection from a quantitative to a qualitative certification standard. We explain the importance of ensuring certification reports are complete and public for the purposes of independent verification. We recommend requiring the use of zero-knowledge proofs ("ZKPs," a cryptographic technology that allows users to prove that they meet an age threshold without revealing their actual age or identity) to support privacy-preserving age assurance once ZKPs are widely supported by mobile operating system vendors. Finally, we suggest other changes to the privacy, security, and parental consent components to make them implementable.

Our comments are organized as follows:

- Section II provides background on measuring age assurance success and failure rates,
- Section III addresses accuracy minimums,
- Section IV discusses certification requirements,
- Section V focuses on privacy, security, and zero-knowledge proofs, and
- Section VI discusses parental consent.

II. Background: Age Assurance Success and Failure Rates

There are a wide variety of different age assurance methods, including those based on commercial and government records checks, (facial) age estimation, government-issued IDs, and behavioral signals.[3] No age assurance method is capable of providing accurate results for all users. Age estimation techniques inherently have a margin of error and thus cannot accurately distinguish between adults and minors close to the age threshold. Age inference methods return unknown results for users with an insufficient profile in the records being consulted. Methods based on proof of identity are unable to verify the adult status of users without government-issued IDs. Because no single method is sufficient to provide a practical age assurance system, covered operators of online platforms need to offer a variety of methods.

A viable age assurance system should meet two accuracy goals:

- Effectively exclude users under the threshold age (minimize the false accept rate).[4]
- Effectively allow adults to demonstrate that they are above the threshold age (minimize the false reject rate).

[3] We rely on the definitions of "age assurance," "age estimation," and "age inference" provided in section 700.1 of the proposed rules throughout this section. See New York Office of the Attorney General, "Notice of Proposed Rulemaking."
[4] The proposed rule uses the terms "false negative" for incorrectly identifying an adult as a minor and "false positive" for incorrectly identifying a minor as an adult. See New York Office of the Attorney General, "Notice of Proposed Rulemaking." This is confusing in that one can think of the purpose of age assurance as classifying age accurately, regardless of whether a minor or an adult is attempting to prove their age. For clarity, we use the terms "false reject" and "false accept" in this comment and would recommend that the rule be adjusted to do the same.

There is an inherent tension between these goals, and between the *accuracy* of an age assurance system and the *availability* of the online platform (or certain features of the platform) to the user population. The more stringent requirements are applied for adults to prove their eligibility (thus reducing the false accept rate), the greater the number of adults who will be incorrectly excluded (increasing the false reject rate).

A common source of confusion is to fail to distinguish between two definitions of "success" in an age assurance method: returning any answer versus returning a *correct* answer. For example, if a facial age estimation system counts every attempt as a success when the image provided by the user is of high enough quality that the system can issue an estimated age, the system operator may be able to claim that it achieves a high success rate. In this case, the rate of failure only accounts for what the proposed rule calls "inconclusive age assurance outcomes," in which the system provides no age estimate. This conflates falsely rejected attempts (where adult users' ages are estimated to be below 18) with successes.

In an age assurance system designed to maximize accurate age estimation, it is important for the "success" condition to include only those instances when the system delivers correct answers. Whether the system returns an inconclusive age assurance outcome or a false reject, the user should have an opportunity to try again with a secondary age assurance method provided by the online platform. With the exception of the situation where a minor presents a government-issued ID for age-proofing, age assurance methods cannot provide a definitive answer that a user is under 18. In most common use cases arising under the regulations, we do not expect minors to present identification to prove that they are under 18, because they are likely to be defaulted into the minor-specific experience anyway. As a result, users who are rejected by one method should not be rejected entirely, but instead need to be offered the opportunity to demonstrate their age via another method. The overall failure rate of an age assurance system is defined by the number of users who are incorrectly classified, whether that incorrect classification is a result of an inconclusive age assurance outcome or of incorrect results.

III. Accuracy Minimums

The proposed rule requires that operators use age assurance methods that conform to two separate accuracy standards:

- An "accuracy minimum" which sets an upper limit on the false accept rate (section 700.1(b)).
- A "total accuracy minimum" which sets an upper limit on the combined false accept and inconclusive age assurance outcome rate (section 700.1(gg)).

All age assurance methods in use by a covered operator must meet the accuracy minimum and the covered operator must provide at least one age assurance method which meets the total accuracy minimum. Our comments in this section address how the two accuracy minimums are defined, the

handling of inconclusive age assurance outcomes, and the incorporation of a numeric circumvention detection standard. Section D provides proposed changes to the rule text to effectuate our suggestions regarding accuracy minimums.

A. Accuracy minimums must better accommodate availability

Section 700.1(b) defines the "accuracy minimum" as the following (we have substituted "false accept" terminology for clarity):

(1) A false accept rate not exceeding 0.1% of minors ages 0 to 7; 1% of minors ages 8 to 13; 2% of minors ages 14 to 15; 8% of minors age 16; 15% of minors age 17; and
(2) A rate of detecting method circumvention that meets or exceeds 98%.

Section 700.1(gg) defines the "total accuracy minimum" as:

(1) A combined rate of false accepts and inconclusive age assurance outcomes not exceeding 0.1% of minors ages 0 to 7; 1% of minors ages 8 to 13; 2% of minors ages 14 to 15; 8% of minors age 16; 15% of minors age 17; and
(2) A rate of detecting method circumvention that meets or exceeds 98%.

The stated intent of the "total accuracy minimum" is to ensure that "at least one age assurance method offered by a covered operator can be successfully completed by virtually all users while allowing covered operators to adopt a waterfall of methods." But "successful" completion is not the same as accurate completion, nor is it a valid goal in the context of the enforcement of the SAFE for Kids Act. The conditions specified in the definition of "total accuracy minimum" do not ensure that a covered platform offers at least one method that is sufficiently accurate, nor do they ensure that the combination of multiple methods offered would make sufficiently accurate age estimation available to "virtually all" of the user population. The conditions merely establish thresholds for how often an age assurance method must deliver an answer, where an undefined percentage of answers could be false rejects.[5]

The NPRM focuses on three main age assurance methods:

- Age inference based on email address (used for commercial/government records checks)
- Facial age estimation
- Age verification via presentation of government-issued identification

[5] Because the rule requires operators to treat users as minors by default, it is not an issue if minors are unable to "successfully" complete accurate age assurance procedures, as the result will be the same in any case, namely, that they are given the safer default experience intended for minors.

Available evidence indicates that all of these methods will fail to validly identify many 18-29 year olds as being over 18, as shown in the examples in the table below:

Method	**False reject rate for 18-29 year olds**	**Source**
E-mail based age inference	11-25% for Verifymy	Verifymy white paper:[6] 11% for males 18-24, 18% for females 18-24, 25% for males 25-29, 23% for females 25-29.
Facial age estimation	Up to 53% for Yoti	Australia Age Assurance Technology Trial[7] Report: 53% for 18, 50% for 19, 32% for 20, 11% for 21, 0% for 22+.
Government-issued ID	4-15%	2020 American National Election Studies survey-based estimate of the fraction of 18-29 year-olds without a valid government-issued photo ID:[8] 15% for 18-19, 4% for 20-24, 5% for 25-29.

With high false reject rates across these three methods, no single age assurance system would be sufficient to allow "virtually all" adults to accurately demonstrate their age. Many adults would be rejected when trying to prove that they are over 18.

To create requirements that would more effectively achieve the goal that "virtually all" users can accurately prove their age, we recommend revising the rule as follows:

- Require that each of the age assurance methods offered by an operator have a maximum false accept rate as currently specified in 700.1(b)(1).
- Delete the definition of "total accuracy minimum."
- Add a new definition of "collective maximum failure rate" and require that collectively the age assurance methods offered by an operator have a maximum false reject rate. In other words, the rules should establish the maximum percentage of users who are unable to demonstrate their age via any of the methods offered by the covered operator. The Office of the Attorney General could choose the maximum percentage to reflect the "virtually all" objective stated above.[9]

[6] Verifymy, "Innovative age assurance," 18.
[7] Age Check Certification Scheme, "Age Estimation Test Report," Table 2.
[8] Hanmer and Novey, "Who Lacked Photo ID in 2020?", 3.
[9] Given the relative lack of widespread deployment of age assurance on the internet, sufficient empirical evidence does not exist to ground the selection of the maximum false reject percentage.

Section III(A)(18) of the NPRM argues that it is unnecessary to specify a minimum false reject rate because covered operators have an incentive to maximize the number of users who are able to successfully demonstrate their age. However, in practice covered operators need to balance access with the cost and operational complexity of offering multiple methods and supporting users through a waterfall when individual methods do not work for them. It should not be assumed that compliance with the rules will naturally incentivize all covered platforms to make sufficiently accurate age assurance available. The rules should make clear that covered operators have a responsibility to provide sufficiently accurate age assurance to all users who wish to prove that they are adults, rather than rejecting those for whom it is inconvenient or costly to determine their age.

B. Inconclusive outcomes are not clearly defined

In addition to the changes described above, we believe it is necessary to clarify how to handle the inherent uncertainty in age estimation techniques. Internally, these techniques operate by estimating the probability that the subject's true age has a certain value. The sum of the probability distribution across a range of possible age values is then calculated and used to effectuate the age assurance process.

Many systems operate by providing an estimate for age, typically by using the age value within the probability distribution for which the probability of the user's true age matching that value is the highest. However, this is just the *most likely* age of the user. Depending on the age assurance technique used, there may still be a high probability that the user's true age is higher or lower than the most probable estimate. For example, the best facial age estimation algorithms in the National Institute for Standards and Technology (NIST) Face Analysis Technology Evaluation had over 25% chance of estimating a 17 year-old as 18 or over, which would correspond to a false accept rate of over 25%.[10] This is far in excess of the accuracy minimum in the proposed rule.

As a practical matter, age estimates near the threshold need to be treated as indeterminate. Many systems address this by padding age thresholds (e.g., requiring that the most probable age estimate be 21 or above if the true threshold is 18) in order to obtain lower false accept rates,[11] which comes at the cost of higher false reject rates.

The proposed rule is unclear on whether a case where the user's most probable age estimate is above the threshold but not above the padded threshold ought to be treated as inconclusive (thus counting against the total accuracy minimum) or as an age assurance outcome indicating that the user is below the threshold (although the examples suggest that they ought to be treated as the latter).

A better approach would be to treat these outcomes as inconclusive, just as with cases where the user's image is too low quality to produce a result. This approach would incentivize greater availability of sufficiently accurate age assurance because adult users who get falsely rejected would count

[10] Hanacek, "FATE Age Estimation and Verification," Figure 10.
[11] See Age Check Certification Scheme, "Age Assurance Technology Trial."

toward the collective maximum failure rate. Such an approach would be more consistent with the intent that "virtually all" users be able to successfully prove their ages.

C. Circumvention detection should follow a qualitative rather than a quantitative standard

The proposed rule requires that covered platforms be able to detect 98% of circumvention attempts in annual certification testing. Specifically, section 700.5(b)(4) requires:

> detection of method circumvention through testing consistent with a nationally or internationally recognized standard, or if none is available, including a variety of attack vectors weighted to reflect the most prevalent risks, with documentation of the quantity and type of attack methodologies tested

This requirement implies that covered operators would need to test for a broad range of types of potential attacks. However, Section III(E)(2)(b) of the NPRM seems to focus specifically on biometric attacks, referencing ISO/IEC 30107:2023. There are a variety of system-level circumvention techniques that do not require any kind of biometric attack, for example using a virtual private network (VPN) to appear to be accessing an online platform from out of state, or using a device that is configured to circumvent device-based age assurance. Early evidence suggests that some of these techniques may already be prevalent in other states where age assurance is required by law.[12]

Enforcing a cumulative 98% detection threshold against all of an age assurance method's most prevalent circumvention attacks raises significant trade-offs related to privacy and openness when it comes to these system-level circumvention techniques. For example, detecting VPN use may require the covered operator to collect privacy-invasive granular location information from the user's device, and preventing device-based age assurance circumvention may limit which users can access an online platform to only those with compliant devices. Using a single cumulative 98% threshold means that as system-level circumvention techniques become more prevalent, covered platforms may be compelled to take increasingly privacy-invasive steps or to restrict the device ecosystems from which users can access their services to remain compliant.

Furthermore, circumvention rates only make sense in the context of a specific set of circumvention techniques. Because some techniques are more effective than others, any measurement of circumvention success inherently depends on the set of techniques tested and their frequency in the measurement set. Without a standardized reference set of techniques, circumvention rates do not provide much information about the true strength of an age assurance method's defense against circumvention.

For these reasons, we do not believe that establishing a numerical circumvention detection threshold is an effective way to assess covered online platforms' efforts to prevent circumvention, nor does it

[12] Lang et al., "Do Age-Verification Bills Change Search Behavior?"

reasonably balance the need for circumvention resistance with the important values of privacy and openness. Rather than requiring the circumvention detection threshold in the definition of "accuracy minimum" and "total accuracy minimum," we recommend that the OAG add into the certification process a qualitative standard for circumvention detection and mitigation in addition to the testing requirement in 700.5(b)(4), similar to the standard established in 700.7 for data use and protection. This qualitative standard could require covered platforms to take appropriate steps to mitigate the most prevalent circumvention risks. Ofcom has taken a similar qualitative evaluation approach in its guidelines for effective age assurance.[13]

D. Proposed rule text

We suggest the changes below to effectuate the recommendations made throughout Section III.

Section 700.1(b):

> (b) *Accuracy Minimum*. The term *Accuracy Minimum* means:
>
> > (1) a rate of false ~~positives~~ **accepts** for an age assurance method that is equal to or less than the following: 0.1% of minors ages 0 to 7; 1% of minors ages 8 to 13; 2% of minors ages 14 to 15; 8% of minors age 16; 15% of minors age 17, excluding failures or refusals by a user to provide requested data and inconclusive age assurance outcomes**.**~~; and~~
> >
> > ~~(2) a rate of detecting method circumvention for an age assurance method that meets or exceeds 98%.~~

Section 700.1(t):

> (t) *Inconclusive Age Assurance Outcome*. The term *Inconclusive Age Assurance Outcome* means following receipt of all requested information from a user, and absent detection of method circumvention, a determination that the age assurance method cannot provide ~~an age or~~ age status for that user, **including cases where an age estimate is insufficiently certain to definitively determine that a user is a minor.**

Section 700.1(gg): Delete the definition of *Total Accuracy Minimum*, replace it with a new definition of *Collective Maximum Failure Rate*, and replace all instances of *Total Accuracy Minimum* in the rules with *Collective Maximum Failure Rate*.

> ***Collective Maximum Failure Rate*. The term *Collective Maximum Failure Rate* means the percentage of users receiving either a false reject or an inconclusive**

[13] "Identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children and where it is reasonable to assume that children may use them." See Ofcom, "Quick guide to implementing highly effective age assurance."

> **age assurance outcome for all of the age assurance methods offered by the covered operator is equal to or less than [X%].**

Section 700.5: Add a qualitative circumvention detection and mitigation requirement.

> **Covered operators must document in their written reports compiled under subdivision (b) of this section the appropriate steps they take to mitigate the most prevalent method circumvention risks to each age assurance method offered by the covered operator.**

IV. Certification

We are pleased to see the requirement in section 700.5 that age assurance methods be certified by an accredited provider. This is an important provision, both for ensuring that these methods perform according to the accuracy minimums and for providing accountability to an independent third party and to the OAG about the performance of age assurance methods. Our comments below explain why the certification reports need to be made complete and public, and suggest changes to the sample size standard to make the standard technically sound in the context of accuracy measurements.

A. Certification reports should include all data necessary for independent evaluation

The certification report required in 700.5(b) provides a strong basis for characterizing the performance of age assurance methods. We recommend requiring the following additional items:

- The specific configuration that was used to test the age assurance method, such as the challenge age for facial age estimation.
- A precise description of the test data.
- The overall accuracy of the collection of offered methods, where acceptance by any of the methods is considered acceptance. This allows for the determination of the overall effectiveness of the covered operator's system.

In general, the certification reports should include enough information to enable independent replication of the claimed results.

B. Certification reports should be made public

At present, public information about how age assurance methods work is limited, and very few of the public claims made by corporations about the accuracy and effectiveness of age assurance systems can be independently verified. The promulgation of the OAG's rules provide a unique opportunity to bring increased transparency to this significant change in how users access online services.

To foster increased accountability and allow for independent verification of test results, the OAG should require that the certification reports be made public. Not only would this allow for public evaluation of the results, it will create additional incentives for covered platforms to continually invest in robust and privacy-protective age assurance, since their own certification reports would be publicly comparable with their competitors' reports.

C. Sample size standards should be precisely defined for accuracy measurements

Section 700.5 (c)(1) specifies the following requirements for measuring accuracy:

> Sample size calculation must yield reliable and statistically significant results with a high confidence level and low margin of error using as a baseline the population of the State of New York most recently reported by the U.S. Census Bureau.

These requirements are not technically sound, for three reasons. First, statistical significance is a term from hypothesis testing where the intent is to determine whether it is likely that the experimental results are sufficient to reject one hypothesis. It is not meaningful in the context of measuring a single quantity like accuracy rate.

Second, the reference to the existing population of New York is unnecessary. The accuracy of a sampling procedure is dependent solely on the size of the sample and does not depend on the size of the population from which the sample is drawn, provided that the population is larger than the sample. In other words, trying to make each sample (e.g., of 14-15 year olds) representative of the size of the underlying subpopulation in New York has no bearing on the question of whether the sample size is sufficient to measure accuracy for that subpopulation within a particular margin of error at a given confidence level.

Finally, confidence level and margin of error are dual quantities: a margin of error is computed for a given confidence level. As numeric values are provided for the error rates, it would be best to specify values for what the margin of error must be at a specific confidence interval.

To specify an accurate test standard, we recommend rewriting the text as follows:

> **Sample size must be sufficient to provide results for each measurement accurate to within [X%] at the [Y%] confidence level.**

We do not have a strong opinion about what the numeric values for X and Y should be; the OAG needs to determine what is appropriate given the goals of the regulation and the balance of competing interests.

V. Privacy and Security

Privacy and security of user data are critical factors for any age assurance system, many of which inherently require the use of personal data. Ideally, the covered operator would learn no additional information about the user as a result of age assurance beyond what it already knows, other than that the user is over 18 (in the case of an adult user). Zero-knowledge proofs (ZKPs) are an important cryptographic technology that allows users to prove that they meet an age threshold without revealing their actual age or identity. Our comments below explain why zero-knowledge proofs should be required once the underlying technology is widely available, and how to ensure the data deletion rule covers all industry-standard deletion methods.

A. Zero-knowledge proofs should be required once technically feasible

The text in this rule focuses heavily on procedural controls, namely requiring that the operator delete user information after the age assurance process. These are good requirements but are inherently difficult for the user to verify, and they require users to trust that the covered operator is correctly implementing its data protection measures, which does not always happen.[14] Where possible, it is better to use technical controls, which do not require the user to trust the covered operator.

In the context of server-based age assurance, zero-knowledge proofs of age provide the highest level of privacy. Section 700.4(c)(3) allows the use of ZKPs to satisfy the requirement of providing an age assurance method that does not require furnishing government-issued ID. However, this requirement is insufficient to provide a high probability that a user will be able to anonymously prove their age without furnishing government-issued ID. As written, a covered operator could comply with the rules by providing a non-ID-based method with a high false reject rate, and then require an ID-based method as a fallback. Moreover, both age estimation and age inference allow the operator to acquire enough information to learn the user's identity, in many cases without the user presenting government-issued ID. As such, the rules as written do not safeguard the ability for users to anonymously access covered platforms despite the age assurance requirement.

We recommend that the OAG add a rule requiring covered operators to provide a ZKP-based age assurance method once technically feasible, specifically once one or more of the mobile operating system vendors with the highest market penetration in New York provide a ZKP-based age assurance method that implements a publicly available ZKP technical specification and age assurance protocol. This can be effectuated by adding the following provision to section 700.4(b):

> (b) To determine that a covered user is not a covered minor, covered operators must:
>
> …

[14] Age Check Certification Scheme, "Age Assurance Technology Trial," 108-109; Chia, "ID photos of 70,000 users may have been leaked, Discord says."

> **(2) Provide a zero-knowledge proof age assurance method once technically feasible, specifically once one or more of the two mobile operating system vendors with the highest market penetration in New York provides a ZKP-based age assurance method that implements a publicly available ZKP technical specification and age assurance protocol.**

B. Data deletion should allow for all industry-standard removal methods

Section 700.1(p) defines "delete" as "permanently destroy, remove, or de-identify information." A common practice for deletion is to encrypt the target data record (e.g., a user's information) with a user-specific encryption key and then delete that encryption key while retaining the encrypted data. This allows the operator to avoid having to rewrite what might be years of records.

This common technique arguably does not fall into any of the methods listed in the definition of "delete," but it should. We suggest changing the definition of "delete" to include "make permanently unreadable," as follows:

> *Delete*. The term Delete means to permanently destroy, remove,~~ or~~ de-identify**, or make permanently unreadable** information using reasonable measures to protect against the unauthorized access or use of such information and to ensure that such information may not be retrieved after the deletion process has been completed.

VI. Parental Consent

While the text around age assurance of the user is quite clear, the text around parental consent is less so. Section 700.2 (e)(5) states:

> Methods of verifiable parental consent. Any verifiable parental consent method must:
>
> …
>
> (iii) be reasonably calculated, in light of available technology, to ensure that the individual providing consent is a parent of the covered minor

This is difficult to operationalize, especially in light of the statement in NPRM Section III(C)(4)(iii) that the OAG does not deem the existing methods approved by the FTC under the COPPA rule to be automatically satisfactory. We note that none of those eight methods establishes the parental relationship to the child, but rather they establish that an adult is willing to claim to be the parent. We recommend that the OAG either provide clearer guidance as to what constitutes acceptable methods for verifying parental consent or provide a process for covered operators to determine that their proposed methods are acceptable.

We additionally note that NPRM Section III(A)(28) states that minors acting in the role of a parent can be recognized as parents for the purposes of the act:

> Finally, existing New York State law recognizes that in cases where an individual who is under the age of eighteen also has the status of a parent (by birth or marriage) to a child, that individual is capable of overseeing their child's upbringing. For example, under Pub. Health Law § 2504(2), such individuals may grant consent for the provision of healthcare services to their child. Consistent with these laws and in these limited circumstances, the proposed rule would allow such individuals to be recognized as parents for purposes of the Act.

This claim is inconsistent with the text of the rules. Section 700.2 (e)(5)(i) and section 700.3(e)(i) first require establishing the parent's age status, which would seem to preclude minors acting in a parental role. This speaks to the need for the rules to provide clearer standards for what constitutes parental consent.

Bibliography

Age Check Certification Scheme. "Age Assurance Technology Trial: Part C Age Verification." Australia Department of Infrastructure, Transport, Regional Development, Communications, Sports and the Arts, August 2025. https://www.infrastructure.gov.au/sites/default/files/documents/aatt_part_c_digital.pdf.

———. "Age Estimation Test Report: YOTI." Australia Department of Infrastructure, Transport, Regional Development, Communications, Sports and the Arts, August 2025. https://ageassurance.com.au/wp-content/uploads/2025/08/IndividualTestReport-YOTI_AE.pdf.

Chia, Osmond. "ID photos of 70,000 users may have been leaked, Discord says." *BBC*, October 9, 2025. https://www.bbc.com/news/articles/c8jmzd972leo.

Hanacek, Natasha. "Face Analysis Technology Evaluation (FATE) Age Estimation and Verification." National Institutes of Standards and Technology, August 29, 2025. https://pages.nist.gov/frvt/html/frvt_age_estimation.html.

Hanaoka, Kayee, Mei Ngan, Joyce Yang, et al. "Face Analysis Technology Evaluation: Age Estimation and Verification." https://pages.nist.gov/frvt/reports/aev/fate_aev_report.pdf.

Hanmer, Michael J., and Samuel B. Novey. "Who Lacked Photo ID in 2020?" Center for Democracy and Civic Engagement, March 13, 2023. https://www.voteriders.org/wp-content/uploads/2023/04/CDCE_VoteRiders_ANES2020Report_Spring2023.pdf.

Lang, David, Benjamin Listyg, Bennah V. Ross, Anna V. Musqera, and Zeve Sanderson. "Do Age-Verification Bills Change Search Behavior? A Pre-Registered Synthetic Control Multiverse." Open Society Foundation, March 2025. https://osf.io/vp9z6/files/z83ev.

Moehring, Alex, Alissa Cooper, Arvind Narayanan, et al. "Better Feeds: Algorithms That Put People First." Knight-Georgetown Institute, March 4, 2025. https://kgi.georgetown.edu/research-and-commentary/better-feeds/.

New York Office of the Attorney General. "Stop Addictive Feeds Exploitation (SAFE) for Kids Act: Notice of Proposed Rulemaking." New York State Register, Steptember 15, 2025. https://ag.ny.gov/sites/default/files/regulatory-documents/safe-for-kids-act-nprm.pdf.

Ofcom. "Quick guide to implementing highly effective age assurance." January 16, 2025. https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/age-assurance?url=https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/online-pornography&data=05%7c02%7cJohn.Eccleston@ofcom.org.uk%7c4682c653674e47cace2b08ddacdcf2da%7c0af648de310c40688ae4f9418bae24cc%7c0%7c0%7c638856787157089173%7cUnknown%7cTWFpbGZsb3d8eyJFbXB0eU1hcGkiOnRydWUsIlYiOiIwLjAuMDAwMCIsIlAiOiJXaW4zMiIsIkFOIjoiTWFpbCIsIldUIjoyfQ%3d%3d%7c0%7c%7c%7c&sdata=oTw%2bPsPihVYIZdR4lNwQ311vnlF9fAoonRN6WEJ1qkA%3d&reserved=0.

Verifymy. "Innovative age assurance: Email address as the new benchmark for frictionless age estimation." June 2024.

https://verifymy.io/wp-content/uploads/2024/11/Verifymy-White-Paper-Innovative-age-assurance-Email-address-as-the-new-benchmark-for-frictionless-age-estimation.pdf.

Yoti. "Yoti Facial Age Estimation." July 2025. https://www.yoti.com/wp-content/uploads/2025/08/Yoti-Age-Estimation-White-Paper-July-2025-PUBLIC-v1.pdf.