NYU | CENTER ON TECH POLICY

Policymakers' Guide to Developing **Age Assurance Legislation**

Scott Babwah Brennen, Lama Mohammed, and Afnan Abbassi

After years of debating if and how we should verify the ages of users to restrict minors from accessing harmful digital content, it is happening in the U.S. In the last several years, dozens of U.S. states have enacted laws requiring adult content websites, social media platforms, and app stores to verify users' age. And while we should continue to interrogate the value of requiring age assurance, it is essential that we now also consider how to enact these policies in ways that best protect user privacy, security, and speech rights.

There is no perfect way to implement age assurance; each approach will involve significant trade-offs. For example, while requiring users to submit government IDs may make systems more accurate, it also raises the risk of privacy violations and unjust exclusion. Lawmakers must understand the trade-offs of each option to balance the choices they must make when drafting age assurance legislation.

Below, we offer a guide for state lawmakers considering new age assurance requirements. We map out **six key decision points** that lawmakers will confront in crafting age assurance legislation, elucidating both the set of choices and the trade-offs inherent in each. We draw heavily on what states have already done, or at least considered, to map out real options for state lawmakers.

Accepting that there is no perfect option that can protect user safety, privacy, and rights, our goal is to help lawmakers understand the implications of each choice so they can better balance advantages and disadvantages for their own constituents.

Decision Point 1:
What is age assurance meant to accomplish?

Lawmakers establish age assurance requirements to protect minors. There are several distinct ways lawmakers generally pursue this.

1. **Restricting minors from accessing content or features considered harmful or inappropriate.**

   Trade-offs
   There is broad public support for limiting minors' access to harmful or inappropriate content or features. However, the courts have long held that minors do have limited First Amendment rights. In that there is no single understanding of what constitutes "harmful" or "inappropriate," overly broad definitions may restrict important or beneficial content. Information about sexual health, LGBTQ+ content, or sexually explicit art or literature could be limited by overbroad efforts to limit "inappropriate" sexual content. Similarly, efforts to protect children from "harmful" social media features or platforms likely limit their ability to express themselves. Once they turn 18, minors will gain full access to all legal content. Restrictions on access to social media or AI platforms may leave minors unprepared to use them safely and responsibly.

   Simultaneously, it is impossible to restrict minors' access to content without also imposing some burden on adults. Until *Paxton v. FSC*, the courts had determined that online age verification imposes a significant and unconstitutional burden on adults' access to legal content.

   Finally, broad restrictions on minors' access to content restrict parents' control over what their children access online. Rather than leaving the decision to parents or guardians about what content is appropriate for their children, these approaches relocate that decision making authority to the government.

   Example
   Wyoming HB 43

2. **Requiring parental consent for minors to access apps or content.** Once determined that an account belongs to a minor, some laws require that the account be associated with a parent or guardian.

   Trade-offs

This approach grants parents control over what content their children access. Yet, requiring parental consent poses significant functional challenges. Consent requirements may fail if parents are not actively involved in their children's lives. The burden of repeated consent requirements, such as approving account creations or certain followers, may prove untenable for parents who choose to bypass the system.

Parental consent requirements also pose risks to minors' privacy and autonomy. Giving parents control over how minors express themselves may undercut minors' expressive rights. Notably, parental consent requirements likely pose a greater risk to marginalized and vulnerable minors. Not every minor has a positive relationship with their parents/guardians, nor have the best interests for their child. Parents may not support certain opinions or lifestyle choices.

Example
Nebraska LB 383

3. **Restricting access to certain features or setting defaults.** Rather than limit access to entire platforms, some lawmakers are considering laws that would require platforms to restrict only certain content or features for minors, or to set default settings for minors.

Trade-offs
This approach preserves minors' access to apps while reducing access to the most concerning content or features. It permits apps to target controls to different age groups and has a higher chance of surviving First Amendment scrutiny as narrowly tailored.

However, implementing this poses significant technical challenges. Smaller companies may struggle to accurately age-gate specific features or content. It also poses significant privacy risks. Apps must maintain age data or regularly reassess user ages. This may mean holding on to sensitive data, rerunning age assurance multiple times.

Example
Nebraska LB 504

Decision Point 2:
Who will assess ages?

Whatever the reason for age assurance, legislators must verify which party should be responsible for verifying a user's age; these include:

1. **The user.** Over the past several decades, self-declaration, in which the user attests to being over (or under) an age cut-off, has been the most common method of age assurance.

   Trade-offs
   Self-declaration protects user privacy by preventing apps from collecting sensitive or personal data. In contrast to other approaches, it is both easy to implement and presents little friction for users, most of whom have experience with these systems.

   Self-delaration, however, is easily bypassed. There is nothing preventing minors from lying about their ages, making these systems inaccurate. Increasingly, courts see self-declaration as insufficient protection.

   Example: California AB 1043

2. **First-party service** (The party that grants access).

   Trade-offs
   Targeted applicability. Only apps that pose specific risks to children are required to implement age assurance measures, thereby limiting restrictions on access to legal adult speech. Rather than granting a single company (such as an app store) the ability to control access to thousands of apps, requiring each app that gives access to content to implement age assurance, decentralizes decision making across the entire ecosystem.

   On the other hand, every app that conducts its own age assurance creates significant friction for users. User experience suffers when users must repeatedly undergo age assurance processes. This also likely imposes non-trivial compliance costs—especially for start-ups or smaller companies, which may give larger, more established companies a competitive advantage.

   Similarly, first-party assurance poses significant privacy risks, as every covered app must collect users' personal data, or hire a commercial firm to do so. This vastly increases the number of companies that must collect, store, and process personal data.

   Finally, distributing compliance across the entire ecosystem may lead to reduced oversight and lower compliance rates. Few state Attorney General offices will have sufficient resources to ensure that every app complies. If the likelihood of enforcement is

low, apps may choose not to comply with a specific state's requirements rather than pay the costs of compliance. Alternatively, they may choose not to offer their apps to users in states with these requirements.

Example
South Dakota HB 1053

3. **Intermediary**. Rather than age assurance occurring at the app level, lawmakers can require intermediaries, including app stores, device makers, and operating system developers, to verify users' ages.

   Trade-offs
   Requiring an intermediary to conduct age verification likely lowers the burden placed on users. Rather than having to submit information to every app, they only have to submit information once. Similarly, it reduces the number of companies that must collect personal information, limiting privacy risks. Because the state government only has to monitor a handful of companies, compliance with new laws is likely far higher.

   However, intermediary verification centralizes app access. As a result, a handful of app stores or other intermediaries must decide access to an entire ecosystem of apps. Intermediary verification is a blunt tool. App stores or other intermediaries may have a limited understanding of particular apps, and therefore will have limited ability to make context-specific decisions.

   Finally, intermediary verification has accuracy issues: users often share app store accounts and devices. For example, children often share devices with an adult. Dealing with shared devices and accounts may require ongoing verification rather than a single, one-time check, potentially making this verification method alone insufficient.

   Example
   Texas SB 2420

4. **Independent commercial thirty-party.** If required to conduct age assurance, many online apps would contract with a third-party vendor. A handful of state proposals go further and require that apps employ a third-party verifier.

   Trade-offs
   Third-party vendors help reduce user burden by allowing them to verify their age once rather than separately for each app. Additionally, contracting with a third-party for age verification may lower compliance costs for smaller platforms that lack the funds and infrastructure to build custom systems.

   Nonetheless, this approach risks centralizing sensitive user data in a small number of

companies, making them sought-after targets for data breaches and attacks. The consolidation of personal data to a few companies may also push out smaller age verification vendors, reducing innovation and potential competition.

Example
Ohio HB 96

Decision point 3:
How is liability structured for age-assurance errors?

Liability determines who can legally be held responsible and potentially face fines, lawsuits, or other penalties when age assurance fails, and minors gain access to restricted content or services. While legislation may require certain actors or methods to complete age assurance, questions concerning liability remain: when does it arise, who is liable for errors, and how can that liability be reduced or mitigated?

Who holds liability for age assurance? Once users' ages are assessed or verified, who is liable if there are errors?

1. **Access-based liability (platform/app liable).** This is by far the most common approach, in which the company that grants or denies access ultimately bears the liability for errors. For example, if a 14-year-old accesses restricted content on a social media platform, the platform faces penalties, even if the verification was done by a third-party that incorrectly verified the user's age.

   Trade-offs
   This method offers clear accountability, acknowledging that what matters most is whether or not kids are given access to problematic content. On the other hand, in many cases, verification is done by a third-party company. Many of these third-party verifiers offer limited transparency or oversight into their systems and processes. Because access providers must trust third parties to be accurate and reliable, there will be an incentive to choose larger, more established companies. Over time, this may lead to the consolidation of verification services.

   Example
   Arizona HB 2112

2. **Knowledge-based liability.** Here, an app assumes liability when it *knows* a user is a minor, rather than merely granting access. This is the current framework under the Child Online Privacy Protection Act (COPPA) for companies with "actual knowledge" that a user is under 13, which imposes certain responsibilities. Under COPPA, platforms are not required to verify all ages upfront, but they must act appropriately once discovering a user is a minor.

   Trade-offs
   A framework in which companies do not have a responsibility to proactively verify the age of all users has a long legal precedent, with "actual knowledge" standards existing in many areas. Importantly, this is far less likely to chill legal adult speech, as adults will not have to verify their ages in many cases.

On the other hand, platforms are rewarded for strategic ignorance and have a strong incentive not to determine users' ages. That said, the incentive means *less* personal data is collected and stored, so while this undermines safety-by-design efforts, it remains more protective of privacy.

Knowledge and ignorance can be difficult to define or prove. This means that knowledge standards may result in legal uncertainty. Again, there is a strong incentive for smaller companies to avoid lawsuits, leading to a split in which smaller platforms overcomply, and larger ones may undercomply.

Example
California AB 1043

3. **Process-based liability (verifier liable).** Here, the liability would fall on the third-party verifier hired by an app or platform. We are not aware of any proposals that include this model. However, this is an approach seen in some other high-risk areas, such as online gambling or online alcohol sales.

   Trade-offs
   While the content provider retains liability for giving access to minors, the verifier would also hold some liability for errors. Age verifiers often provide little transparency into their process; this would ensure that vendors are held accountable for their verification process. Importantly, a vendor has no control over how an app uses the information they provide. Likely, this means their liability would be limited to errors in the actual verification of users, rather than the decisions regarding access.

   Example
   Texas SB 650

Liability mitigation/safe harbors. Irrespective of how lawmakers apportion liability, some frameworks include ways for companies to mitigate liability.

1. **Time-based safe harbor (cure periods).** Some frameworks allow companies a certain amount of time to correct their systems once they are informed that they are out of compliance.

   Trade-offs
   This approach prioritizes remediation over punishment, while helping ensure companies are not penalized for mistakes or oversights. However, given the limits of government enforcement efforts, this may mean some companies decide there's little reason to comply until they receive notice they are out of compliance. Regarding age assurance, this means that even after a law is enacted, some children will be exposed to problematic content or features before a violation is cured.

Example
Utah SB 142

2. **Compliance-based safe harbor**. Some laws stipulate that as long as an actor follows a set of proscribed methods or relies on approved providers, they will be protected from liability.

   Trade-offs
   These laws specify the steps companies must complete to comply with the law. Regulatory clarity not only simplifies enforcement but also likely leads to faster adoption, as companies do not want to be notified of noncompliance. At the same time, it acknowledges that age assurance is never perfect, and some minors will slip through. As long as companies implement proper processes, they will not be penalized for failures.

   A "one-size-fits-all" approach, where every company must follow a very specific set of steps, may not work for every company or every situation. Having broader, less tailored laws may lead to overcompliance, in which companies may restrict adults' access to legal content. In practice, compliance-based safe harbors are likely to encourage apps to contract with third-party vendors, leading to a centralization of age assurance in a handful of vendors.

   Example
   Texas SB 2420

3. **No safe harbor.** Rather than permit companies to correct issues or demonstrate a good-faith effort to comply, some laws impose penalties for any failure to comply with the law.

   Trade-offs
   This strict enforcement regime increases the likelihood of overcompliance, as covered companies face severe penalties for any errors. However, age assurance is never perfect, and here companies face consequences even when acting in good faith. This provides a strong incentive to implement strict, intrusive verification methods**.** As a result, adults are likely to be restricted from accessing legal speech.

   Example
   Arizona HB 2112

Decision point 4:

How is age assessed?

There are many ways to assess or estimate a user's age. Each has advantages and disadvantages. We catalogue methods into five groups. Notably, most laws permit multiple

forms of assurance.

1. **Self-declaration.** As discussed above, self-declaration has been the most common form of age assurance.

   Trade-offs
   Self-declaration protects user privacy by preventing apps from collecting personal data. In contrast to other approaches, it is both easy to implement and causes little friction for users, most of whom are already familiar with these systems.

   Self-delaration, however, is easily bypassed. There is little preventing minors from lying about their ages, making these systems inaccurate. Increasingly, courts see self-declaration as insufficient protection.
   Example
   California AB 1043

2. **Government-issued ID.**

   Trade-offs
   Requiring government-issued IDs is a straightforward and accurate method to verify user ages. It produces a clear audit trail that makes it easy for the verifier to defend its decisions.

   However, requiring government IDs will exclude users without current government IDs—disproportionally impacting historically marginalized populations. There is no legal requirement in the U.S. to have a government ID, yet this approach makes it necessary to access online content.

   Some may not feel comfortable submitting government IDs to access online platforms. Requiring government IDs may drive away adults who would otherwise be legally permitted to enter. To circumvent these risks, both adults and minors may use VPNs to mask their location or visit other noncompliant adult platforms.

   Finally, requiring government IDs requires collecting, storing, and processing sensitive data. This is not only expensive and technically complex, but also poses significant privacy risks, as collections of government IDs make them attractive targets for bad actors.

   Example
   North Dakota HB 2380

3. **Commercial third-party data**. Another option is to use existing third-party commercial data sets, such as credit card or mobile phone data, to verify user ages.

Trade-offs
Using third-party datasets is often easy to implement given existing technical infrastructure. Moreover, because these companies already have extensive data on many users, most users will not need to upload new documentation.

There is likely to be little to no transparency into these systems, making auditing extremely difficult. Moreover, while third-party companies hold much personal data, it is not universal, and this framework likely excludes users who do not use credit cards or mobile phones.

Example South Dakota HB 1053

4. **Social attestation.** Some social media apps, including Instagram and Facebook, have permitted one's friends or social relations to attest that a user is an adult. We are not aware of any laws that explicitly permit social attestation as a method of assurance.

   Trade-offs
   These systems do not require verifiers to collect personal data, and do not require users to have government IDs, credit cards, or mobile plans.

   However, social attestation can be slow, complicated to enact, and inaccurate. Users must wait while several contacts verify their age. Platforms have generally limited the number of people a single user can vouch for and must ensure that those vouching for others are themselves adults. Unlike submitting government IDs, this system is relatively easy to bypass. For these reasons, social attestation can be challenging to scale and is likely more effective when only a small number of users need their ages verified.

   Example
   Meta

5. **Biometric age estimation.** Companies can use AI models to estimate a user's age from an image of their face or other biometric markers.

   Trade-offs
   Biometric age estimation offers an accurate, low-friction method of age estimation. Biometric systems can work quickly and can be integrated on platforms in ways that minimize disruption for users.

   That being said, in conducting biometric age estimation, companies must collect and process biometric data. Some have deep concerns about such data collection. At the same time, there is extensive evidence that facial age estimation can be biased against certain groups—in particular, those with dark skin. Even slight differences in accuracy between demographic groups can result in meaningful discrepancies when age estimation must distinguish between a 17-year-old and an 18-year-old.

Example
Virginia SB 854

Decision point 5:
How is age assurance enforced?

Lawmakers must decide *how* age assurance laws will be enforced. Broadly, there are three common approaches.

1. **Public enforcement.** Under existing age assurance laws, the most common approach is to delegate enforcement to the state Attorney General. In some cases, other government agents are delegated enforcement. There are a handful of tools that AGs can bring to bear, depending on how laws are crafted, including civil penalties, criminal penalties, injunctions, settlements, and consent decrees.

   Trade-offs
   Delegating enforcement to the AG or other government offices helps ensure that enforcement targets the worst offenders. Public enforcement can serve as a powerful deterrent, especially when enforcement threatens severe monetary fines, criminal liability, or injunctions. Public enforcement also often leverages existing state resources, meaning states do not need to establish new regulators, a process that usually takes time and money.

   At the same time, most public offices across states are resource-constrained. If the law requires every app to comply, it may be beyond the capacity of public officials to either monitor compliance or bring cases. This means compliance will likely be incomplete, and some offenders—especially smaller players—may gamble that the law will not be enforced. Furthermore, Attorneys General are political figures, meaning in some cases, enforcement may be politicized.

   Example
   South Dakota HB 1053

2. **Private enforcement.** Also called a "private right of action" (PRA), this approach permits users the right to sue a company for violations if they have been harmed.

   Trade-offs
   PRAs provide a strong incentive for companies to comply; lawsuits can be extremely costly. Depending on how damages are awarded, adverse decisions could be substantial. PRAs give users who have been directly harmed a means to seek redress.

   However, PRAs often lead to overcompliance. Companies, big and small, must spend time and money dealing with suits that may be frivolous. Even suits that are quickly thrown out impose costs and burdens on companies. While large companies have in-house legal teams, smaller companies may not and may have to spend significant amounts of money dealing with frivolous lawsuits.

Example
Utah SB 142

3. **Delegated enforcement**. Some laws require intermediaries, such as app stores, to ensure that apps adequately verify users' ages. While public enforcement may be used in some cases, these laws effectively require intermediaries to serve as legal enforcers or risk serious penalties.

   Trade-offs
   Delegated enforcement reduces the enforcement burden on state governments. Intermediaries often have a far better understanding of apps and can assess risks more accurately.

   On the other hand, there will be a strong incentive for over-compliance. Platforms may have little interest in protecting expression rights—likely prioritizing user experience and minimizing litigation risk. This could also lead to significant speech restrictions.

   Example
   Louisiana HB 570

Decision point 6:

How is age assurance data used and retained?

Lawmakers must establish data privacy and usage restrictions, specifically: what signal is produced and shared, how long data is kept, and (how) can it be used for other purposes?

What data is held or output? If the verifier is not the same as the company providing access, there are several options for what age signal the verifier sends to the relying party,spanning from minimal to maximal personal data.

1. **Eligibility signal**. Rather than send personal data, the verifier can share a signal indicating whether the user is over or under 18 (or another age cut-off).

   Trade-offs
   In this approach, the verifier is the only party that has access to users' sensitive or personal data. There is little risk of leaks or identity theft during transmission. The simple thumbs-up or down signal means the verifier can use a token-based system that is highly protective of user identity.

   However, this system likely means that the relying party will be unable to target design or features to specific user ages. Moreover, it may mean that the relying party must regularly request an age check for those under 18. The relying party won't know whether the person is 17 and 10 months old or 13, so they will need to continue checking ages. Finally, these systems are likely to result in a small number of verifiers supplying age tokens to the broader ecosystem. In addition to centralizing access decisions, this can present attractive targets for data theft.

   Example
   EuConsent

2. **Age band or exact age**. Rather than simply a pass/fail signal, the verifier can send either an age band (13-15; 16-17, 18-25, etc) or a user's exact age or date of birth.

   Trade-offs
   This approach limits the personal data shared by the verifier but can provide the relying party with additional data to better adjust content or features. When sharing the birth date, the verifier only needs to send the signal once to the relying party; the relying party can retain the birth date/age and update access as the user ages.

   However, age or birth date counts as personal data in many states, and many relying parties do not want to retain this information about users. Depending on rules around

data reuse (see below), relying parties will be able to use these data to target ads or features.

Example
Utah SB 142

3. **Full identity.** This model is more like a full background check, where the verifier does the legwork, but then shares all the information they collect with the relying party. While legislation usually does not specify that the verifier must send full information to the relying party, some laws do not restrict the types of information that can be shared, leaving this as a legal option.

   Trade-offs
   This is the most transparent approach, permitting the company that gives access—which often bears full liability (see above)—to understand how verifiers make the decisions they do.

   Sharing full identity data magnifies privacy concerns and security risks, as multiple parties may share, process, and store sensitive information. Furthermore, companies may be able to sell or use complete identity information to target ads.

   Example
   Ohio HB 96

Retention. How long is data kept?

1. **Ephemeral.** Some laws require that verifiers or relying parties retain personal information only long enough to complete the immediate age verification task.

   Trade-offs
   While this means that verifiers or relying parties do not hold personal or sensitive information, they must reverify a user's age every time a user wants access to the app—or even to a particular feature. This can impose significant burdens on users and increase risks of data leaks or theft.

   Example
   Missouri 15 CSR 60-18

2. **Time-limited or indefinite.** Verifiers or relying parties can hold age information for a set period of time. In some cases, there may be no restrictions on how long they can hold age data.

   Trade-offs

Permitting verifiers or relying parties to hold verification data or age signals for a period of time means they will not need to ask users to resubmit verification information. However, the longer data is held, the greater the risk of misuse or security breaches.

Example
Tennessee HB 1614

Reuse. (How) can data be reused?

1. **Single-purpose.** Often addressed in "data minimalization" provisions, this framework permits either the verifier or the relying party to use collected data only for the stated verification purpose. Importantly, these restrictions may be included in age assurance laws or may be included in broader data privacy laws.

   Trade-offs
   This is the most privacy-protective framework. It is simple and straightforward, and likely politically popular. It assures users that age verification isn't a pretext for data collection or surveillance.

   That said, depending on how the provisions are written, they may require verifiers to repeat verification. Moreover, without proper exemptions, single-purpose restrictions may prohibit benign or beneficial uses, such as identifying fraud, bots, or other bad actors, completing academic research, or improving product design.

   Example
   Nebraska LB 504

2. **Restricted or open use.** In some frameworks, lawmakers permit companies to use age assurance data for a set of approved uses. While we are not aware of any proposals that explicitly grant open use, if laws do not explicitly restrict secondary uses, they create de facto open use.

   Trade-offs
   There are ways that platforms can use user data to design better products, better identify bots or inauthentic behavior, or better protect users from scams or other hazards. Permitting some secondary use allows companies some latitude to repurpose data to build safer or more efficient systems, while preventing those that pose greater user risk or are likely to be distasteful to users. More permissive data policies can also offer companies space to find new ways to monetize. Not only does this increase companies' incentives to implement age assurance, but it can also help build a stronger, more robust ecosystem.

   Yet, writing fine-grained policies that specify specific permitted uses can be extremely difficult. There are often ways to work around restrictions. Lawmakers may not have a deep understanding of companies' data practices; companies are almost always better

at finding loopholes than lawmakers are at closing them. Open-use of age assurance data means users must trade knowledge about or control over their data use to access legal content. At a time when many are already deeply concerned about platform data practices that condition access to legal content on the collection of personal data and allow companies to do whatever they want with it, raising serious data privacy concerns. Here, it is likely that age verification data will feed into both private and public profiling and surveillance infrastructures. Increased surveillance by both companies and the government can lead to severe speech limitations.

Example
Ohio HB 96