Institute for Family Studies
Strong Families. Sustainable Societies.

Ethics and Public Policy Center

AGE VERIFICATION

Policy Ideas for States

John Ehrett

Clare Morell

June 13, 2023

ABOUT THE AUTHORS

John Ehrett is Chief Counsel to U.S. Senator Josh Hawley on the Senate Judiciary Committee, and he serves as lead Republican counsel on the Subcommittee for Privacy, Technology, and the Law. His work has appeared in a wide variety of scholarly and professional publications, including the *Yale Journal on Regulation*, *Journal of Law and Politics*, *Information and Communications Technology Law*, *American Affairs*, and *The New Atlantis*. He is a graduate of Patrick Henry College and Yale Law School.

Clare Morell is a senior policy analyst at the Ethics and Public Policy Center, where she directs the Technology and Human Flourishing Project. She worked in the White House Counsel's Office and the Justice Department during the Trump administration.

ACKNOWLEDGEMENTS

The conclusions and recommendations in this legislative brief are those of the authors alone and do not necessarily reflect the opinions of our employers, sponsor, or advisers. The authors wish to acknowledge the guidance and direction of Michael Toscano, who facilitated the brief's production. The authors also wish to acknowledge the help of Jonathan Askonas (The Catholic University of America) and Andrew Spear (Strider) for their review of the brief. This brief was edited by Alysse ElHage.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

Introduction

The evidence is increasingly clear that social media is harming young people.[1] As a result, both states and the federal government are considering strategies to better protect kids online. The effectiveness of any policy solution, whether raising the age of social media or requiring parental consent, among others, will hinge on effective age-verification—the ability to determine if a user is above a certain age threshold in order to keep underage individuals from access. This paper examines that topic for policymakers.

To date, "age gates" commonly used by alcohol and gambling websites to limit access to underage individuals—i.e., drop-down menus directing a user to enter a date of birth, which deny access to underage users—are easy to circumvent.[2] Users need only be willing to lie about their age to gain access. This is also true of the toothless age verification approach used by social media companies like TikTok.[3]

Dissatisfied with such lenient approaches, legislatures have already begun imposing more robust requirements. Most notably, in March 2023, Utah Governor Spencer Cox signed into law S.B. 152, a bill that would require both age verification for Utah-based social media users and parental consent for users under the age of 18 to open or operate an account.[4] The law directs Utah's Division of Consumer Protection to make the rules to establish the means by which a company may meet the age-verification requirements.

Some technologists and industry-aligned groups have fiercely criticized proposals like S.B. 152. The Electronic Frontier Foundation alleges that "[a]ge verification systems are surveillance systems," which "would lead us further towards an internet where our private data is collected and sold by default."[5] Raising similar privacy concerns, industry group NetChoice alleges that under an age-verification regime, "[e]veryone would be required to hand over a form of identification to tech companies to verify age and identity every time you access websites."[6]

However, these claims need not be true, as our paper will explain. It is important, though, to acknowledge that any age-verification system will entail tradeoffs. The question is which tradeoffs, in the judgment of elected or appointed policymakers, are worth making. One prominent decision point with age-verification involves deciding where to land between the level of effectiveness of the method and its level of intrusiveness to the individual user.

In developing age-verification legislation or policy, lawmakers will have to formulate answers to two central questions:

(1) What information needs to be collected for age verification purposes?
(2) What entity will be tasked with actually performing the age verification and handling the associated personal data?

The answers to these two questions will necessarily form the core of any age-verification strategy. And each can be answered in a number of different ways, as this memo will discuss. Once a core framework has been decided upon, other relevant considerations, like possible enforcement mechanisms, can then be incorporated.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

1. Types of Age Verification Information

The first question policymakers must answer is straightforward: what kind of information constitutes appropriate proof of user age? This is a perennial problem. For instance, at a point of sale for alcoholic beverages, a driver's license is typically accepted as sufficient evidence that a purchaser is legally eligible to purchase alcohol. But as any bartender is well aware, fake IDs remain a problem, and the challenge is compounded when such age checks are mediated digitally.

Today, the types of information used for age verification generally fall into four major categories. Each potential approach has pros and cons of its own.

A. GOVERNMENT RECORDS

The first type of approach for an age-verification law or other policy is to require users to upload scans or photographs of government-issued identity documentation as a condition for social media access. Such identity documentation might include a driver's license, a passport, a birth certificate, or a state-issued identification card, to name just a few examples.

This approach has several advantages. First, such information is—at least in principle—uniquely capable of verification. Unlike many other forms of identification, license numbers and other identifying information can be cross-referenced against public records databases to verify that a user is in fact who he or she claims to be.[7] Second, this strategy benefits from the deterrent effect created by existing laws prohibiting falsification of government documents.[8] Obtaining a fake ID is not costless or without risks: users tempted to try to manipulate such an age verification system would face a significant, and legally dicey, hurdle.

However, there are also downsides to this approach. Most notably, an ID-based strategy requires disclosure to the verifying entity of a significant amount of personally identifiable information. Submitting a photograph of a driver's license, for instance, may involve disclosure of a unique identifying number, the holder's date of birth, an image of the holder's face, the holder's height and weight, and so forth. There are valid privacy concerns, as uploading an ID to an online platform is not the same as flashing an ID to enter a bar. Information provided online is not "forgotten" in the same way. Accordingly, the risks associated with potential breaches of the verification system are compounded.[9]

There are other concerns also. Users without government-issued identification documents may struggle to access social media platforms requiring such information as a condition of entry.[10] Additionally, it is not necessarily straightforward or costless for entities to check user-submitted information against public databases. And finally, as with any age verification approach based on images of identification documents, it is possible in principle for users to upload false information, which may not be reliably screened out.

B. FINANCIAL DOCUMENTATION OR TRANSACTIONAL DATA

An age-verification law or policy might require users to provide a credit card number or analogous piece of financial information, such as a reversible charge made to a credit card or bank account. Typically, such a number is understood to serve as a proxy for the user's age. Websites in the adult

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

industry have employed this method for years.[11] The major advantage of such an approach is that it does not require users to divulge more sensitive personal information or significant amounts of user data: a credit card number or its equivalent, taken alone, does not identify the user in any particular way. This method has been employed by online gambling websites and alcohol delivery sites.

The principal downside of this approach is that credit card ownership is highly unreliable as a proxy for age: as early as 2002, the equation of credit card possession with adulthood was viewed as outdated.[12] A determined minor can readily use a parent or sympathetic adult's credit card as "proof of age," undermining the point of age verification in the first place.

C. BIOMETRICS

Another approach that an age-verification law or other policy might take is to use biometric markers associated with age, such as an iris or fingerprint scan of the user, or a scan of the user's face, which is then evaluated by an artificial intelligence tool (an approach currently being developed by companies such as FaceTec and Yoti).[13] The principal advantage of a biometrics-based model is that it is uniquely difficult to falsify one's physical attributes, so the reliability of the method is in theory high.

However, there are significant downsides and serious concerns to a biometrics-based approach to age verification. Most notably, privacy concerns are at their zenith here. Entities' collection of intimate information on users' physical attributes, in an environment where such information is commonly used to safeguard users' access to their sensitive financial records, would make data repositories uniquely tempting targets for identity thieves.

Additionally, from a civil liberties perspective, this data would be particularly ripe for seizure by outside parties, such as hostile state actors, or weaponization against individual users. At bottom, user biometric information is extremely sensitive data to be handing over to private technology companies, who have a significant incentive to exploit it for their own business purposes—that is, maximizing profits by maximizing user engagement—or other methods of social control.

Furthermore, depending on the design, a biometrics-based approach might end up conditioning users' access to social media platforms on their purchase of cumbersome additional hardware, such as retinal or fingerprint scanners.[14] This would impose a significant burden on end users. Furthermore, less intrusive approaches, such as facial scans, are easily sidestepped: a user need only hold up a photo of a different person to "prove their age."[15]

D. INFERENTIAL/INDIRECT INFORMATION

An age-verification law or other policy might adopt a "behavioral" approach, assessing the user's internet use patterns and extrapolating from there whether the user is old enough to access a particular site. This might take the form of a review of the user's internet browsing history writ large, or—in the case of a platform operating across multiple online sectors—an analysis of the individual's use of the publisher's own internet products.[16] For instance, a platform operator at Google might conclude that a user who spends the vast majority of his time watching children's television programs on YouTube is probably not old enough to use one of Google's social networking tools, and might deny access on that basis. The advantage of this approach is that no "raw" personal information (e.g., a user's actual date of birth, driver's license number, or similarly

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

identifying information) is shared with the entity performing the age verification; the user remains wholly anonymous.

However, this approach to age verification amounts to little more than strategic guesswork. Additionally, its effectiveness is entirely dependent on the use of "tracking cookies" and other technologies designed to surveil a user across the internet, an approach which presents significant privacy concerns of its own.

2. Mechanisms of Age Verification

Once policymakers have settled on the type or types of acceptable information to be used as the basis for age verification, they must determine the entity tasked with *performing* the age verification. Decisions at this stage will ameliorate or exacerbate tradeoffs associated with the type of data collected. For instance, the privacy concerns associated with a more "intrusive" form of data collection can be mitigated or worsened by assigning different entities to actually carry out the age verification and potentially data storage.

A. FIRST-PARTY VERIFICATION

On this model, the entity tasked with age verification is the same entity that operates the platform for which age verification is required. For instance, if access to Facebook is age-restricted, a first-party verification model would assign Facebook the responsibility of checking users' ages prior to permitting access.[17]

An advantage of a first-party model is that it allows policymakers to more easily hold companies accountable for failures to conduct necessary verification, avoiding the finger-pointing potentially associated with more complex verification mechanisms. To the extent age-verification requirements are deployed as part of a larger regulatory effort directed against social media platforms, this consideration is not marginal.

To be sure, this model imposes substantial compliance costs on platforms, which must independently undertake to verify the ages of all users seeking to access their platform. Depending on one's perspective, this might be either an advantage or a disadvantage. Critics of social media might conclude that as the principal instigators of the apparent crisis, companies should bear the responsibility and costs of age verification. Tech company defenders, in turn, might argue that such a regime is cripplingly expensive and rife with failure points.

A notable disadvantage of this model is that it necessarily places raw identification data in the hands of platforms with a powerful incentive to misuse that information. A company collecting images of users' passports, for instance, would thereby acquire significant additional information on their users' demographic profiles. That information, in turn, could be readily repurposed for advertising purposes or otherwise monetized.

B. THIRD-PARTY VERIFICATION—TRADITIONAL MODEL

On this model, the entity tasked with age verification is *not* the same entity that operates the platform for which age verification is required. In recent years, a large number of age-verification companies, such as Privo, have emerged to offer such third-party services to companies offering internet products aimed at minors. Companies operating age-restricted platforms might contract

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

with age-verification companies to verify user ages, and then permit or deny access to their platforms on the basis of the third-party company's determination of user age.

A third-party verification model could deny first-party companies access to age-verification data collected from their users, reducing the risk that such data might be monetized or otherwise abused. The corollary of this, of course, is that first-party companies could be less easily faulted for breakdowns in the age-verification process.

Additionally, from a privacy perspective, a third-party entity conducting age verification across multiple sites would constitute a single point of failure. A bad actor who breached such a third-party entity's security systems could access personal information associated with users' identities on a substantial number of age-restricted websites, compounding the harm of identity theft.

This is a model that the state of Louisiana has recently adopted to implement the state's new age-verification law for adult websites. The state partners with a third-party company called LA Wallet to provide digital IDs to its residents. The law allows a few methods for age verification, and the one that is likely most protective of user privacy employs a two-step verification process using LA Wallet. The user does not need to upload any identifying information to the site; the company merely provides a user with a key to enter on the site that indicates they satisfy the age requirement. Chris Griswold, policy director at American Compass, has discussed the possibility of the federal government administering such a system on a national level for age-verification:

> *Administratively, the government is best suited to provide verification of the information it creates and maintains—such as the birth dates public entities certify and the Social Security numbers they issue. Public services already put this information to widely accepted, non-controversial use—as with the federal E-Verify system, or the Social Security Administration's (SSA) number-verification service, which employers routinely use to confirm that new hires have a valid identity. It would not constitute an engineering marvel to create a simple public system that allows Americans to use these data to anonymously prove whether they exceed a given age. One possibility would be for the SSA to offer a service through which an American could type his Social Security number into a secure federal website and receive a temporary, anonymized code via email or text, like the dual-authentication methods already in widespread use. Providing this code to an online platform could allow it to confirm instantly with the SSA whether the user exceeds a certain age without further personal data reaching the platform or the government.*[18]

C. THIRD-PARTY VERIFICATION—ZERO KNOWLEDGE PROOF

On this model, age verification is accomplished through the use of a "zero-knowledge proof"—a cryptographic method that "demonstrates how a prover can confirm any particular statement without giving the verifier any vital information or disclosing information related to the witness."[19] William Buchanan, professor of applied cryptography at Edinburgh Napier University, laid out as early as 2020 how a zero-knowledge proof could be deployed in the context of age verification:

> *A typical problem is that Peggy has to prove that she is now 18 and can be served in a bar—let's call the bar "Crypto Nights." So, the bar is very strict on age requirements and Victor guards the door and makes sure that everyone is over 18….*

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

> *Peggy says that she is over 18, and Victor asks for her age. . . . So Peggy asks Victor, "Who would you trust to verify that I am old enough?" "Well, I trust [Trent] to verify!"*
>
> *Now Peggy goes to Trent [and] tells him the age she has to prove, and Trent sends a random seed to her and generates a signed proof statement that she can give to Victor. Peggy now encrypts her age, and passes this to Victor, and with the two elements, Victor is able to let her in. Victor has no idea how old Peggy is.*[20]

In simpler terms, this verification model can be envisioned as a form of third-party verification that does not rely on the prior existence of a relationship between the verifying third-party entity and first-party point of access.

Technologically speaking, the encryption contemplated here relies on a process known as "hashing"—a means by which a computer can (1) obtain an alphanumeric derivative that is mapped to a specific piece of information, and then (2) verify that the underlying content of two pieces of information are the same without allowing that information to be readily reconstituted from the "hash value," or the derivative.[21]

Hashing is often employed in the context of policing child sexual abuse material online. Illegal images—which, as digital files, are ultimately reducible to strings of computer code—have their constitutive code "hashed" by law enforcement or transformed into much shorter alphanumeric derivatives to create a "hash value." The "hash values" of image files uploaded to internet sites are routinely checked against the known "hash values" of illegal images to determine whether an uploaded image does, in fact, depict child abuse. A human operator can then review the images flagged through such a process.[22] The strength of this method is that, at early stages of the investigative process, such "hash checks" do not require exposure of the underlying information.

In the context of age verification, this approach would involve a third-party verifier (perhaps a government agency itself) providing programmatic access to validate a hashed identity provided to a third party. Users wishing to verify their age (or, by the same process, some other characteristic) could navigate to a first-party verification website and use a verification process (perhaps biometrics) to generate a hash. If additional protection of privacy or anonymity is desired, the data according to which the user's age was verified could at this point be deleted, leaving a validated hash in the third-party database. This hash value can then be checked, according to a hashing procedure that has been approved by a third party ("Trent" in the above example), to provide a yes-or-no answer to the question of whether the user is old enough to access the site, without revealing the user's underlying age or any other details about the user.

As should be evident, there are significant privacy advantages to this approach. Most notably, no "raw" data is exposed to first-party entities. Furthermore, the trusted third-party entity involved need have no relationship to the first-party point of access (e.g., the social media company) beyond merely "being trusted." This helps avoid any conflicts of interest associated with an ongoing contractual relationship.

Such a model could be rolled out in partnership with state DMV offices or other state administrative or public records offices: an individual might go in person to a physical location, present their ID or other government documents, and receive a random seed and proof statement to use for an online platform indicating that they are over the age threshold. This protocol can be run at any degree of decentralization—it would even be theoretically possible to run it on the blockchain with distributed

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

biometric access points generating hashes without creating additional centralized databases of user identity information.

The downside of a zero-knowledge proof approach is that it adds an additional layer of complexity to the process, potentially resulting in cumbersome requirements for end users. Additionally, its effectiveness is also contingent on the existence of trusted third-party entities capable of providing users with the necessary encryption/decryption tools. Finally, "raw" user data would still be collected by the third-party entity.

3. Further Legislative Design Considerations

Policymakers' answers to the two core questions—what type of data are accepted, and who stores and uses that data to verify users' ages—will form the underlying structure of any age verification approach. Once that structure is in place, a wide range of additional provisions tailored to achieve specific legislative objectives can be incorporated.

A. AVOIDANCE OF CONTENT-BASED RESTRICTIONS

Any age verification regime, however structured, must be carefully designed to withstand First Amendment scrutiny. Particularly relevant to this discussion is the history of caselaw regarding past congressional attempts to regulate online pornography—attempts which the Supreme Court twice rejected on constitutional grounds.

In 2004, the U.S. Supreme Court decided the case of *Ashcroft v. American Civil Liberties Union*, which dealt with the constitutionality of the Child Online Protection Act (COPA). COPA was a federal law "impos[ing] criminal penalties of a $50,000 fine and six months in prison for the knowing posting, for 'commercial purposes,' of World Wide Web content that is 'harmful to minors.'"[23]

Notwithstanding that requirement, however, COPA provided an affirmative defense to any website proprietor who established age controls for their site, such as "a credit card, debit account, adult access code, or adult personal identification number" or "any other reasonable measures that are feasible under available technology."[24]

COPA represented the second time that federal lawmakers had attempted to establish aggressive child protection measures for the internet. The first effort was codified in age-restriction provisions of the Communications Decency Act of 1996—provisions which the Court had found unconstitutional in the 1997 case of *Reno v. ACLU*, on the grounds that they were "not narrowly tailored to serve a compelling governmental interest and because less restrictive alternatives were available."[25]

The *Ashcroft* Court found that even COPA's more tailored framework still violated the First Amendment, reasoning that "[c]ontent-based prohibitions, enforced by severe criminal penalties, have the constant potential to be a repressive force in the lives and thoughts of a free people."[26] The Court emphasized that "the statute labels all speech that falls within [its] definitions as criminal speech," a designation that raised substantial First Amendment concerns.[27]

To be sure, the Court acknowledged the problem of children's exposure to harmful material online, and telegraphed support for a filter-based approach that would,

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

> *impose selective restrictions on speech at the receiving end, not universal restrictions at the source. Under a filtering regime, adults without children may gain access to speech they have a right to see without having to identify themselves or provide their credit card information.*[28]

In other words, the responsibility for protecting children online lay with parents and individual households, not the proprietors of websites.[29]

Ashcroft is widely taken to be the main constitutional obstacle to legislation imposing age-verification requirements. In particular, the Court's remarks on user-selected filters as preferable to "back-end" access restrictions may imply hostility to *any* type of age-verification legislation. However, the remainder of the Court's remarks on the subject should be noted: "Above all, promoting the use of filters *does not condemn as criminal any category of speech*, and so the potential chilling effect is eliminated, or at least much diminished."[30] In other words, the central constitutional problem with COPA appears to have been its classification of some forms of speech (obscenity in this case) as uniquely harmful, and then the use of such a classification as the basis for functionally requiring the creation of age-verification systems.

Age-verification regimes need not make this mistake. Left open in *Ashcroft* is the possibility of regulating modes of communications technology *as modes of technology*, irrespective of whatever speech happens to be disseminated therein. Restricting minors' access to certain types of communication platforms does not *necessarily* imply a judgment about the content of the speech on the platform. Rather, it can imply the legislative determination that the design decisions associated with social media—such as perpetual notifications, infinite scroll, the deliberate amplification of content, use of algorithms to recommend content, and so forth—are themselves causing harm. For a platform to argue that such design decisions *in themselves* are subject to First Amendment protections, as somehow "expressive," would stretch existing First Amendment doctrine to the breaking point, making virtually all product regulation impossible. Courts would likely not favor such claims.

Accordingly, legislators should take pains to ensure that they are in fact regulating the instrumentalities of online communication *as instrumentalities*, rather than as repositories of particular categories of user speech. To that end, policymakers should avoid including legislative findings denouncing the prevalence of "harmful content" on social media, which a reviewing court would likely identify as evidence of legislative purpose.

B. PARENTAL CONSENT AND DATA ACCESS

As part of a larger age-verification strategy, policymakers may wish to include provisions requiring that a minor user's parent or guardian consent to the formation of a minor's account and/or be notified of the minor's account creation, and/or given access to and control over the account. Parental-access provisions add an additional dimension to the problems of data collection and validation. The entities involved must establish not just that the age of the primary end user associated with the account is above the age threshold but also must verify the age of the parent or guardian, as well as the existence of a custodial relationship between that user and their ostensible parent. In principle, however, this is simply another category of information that must be provided in some fashion and then verified according to one of the methods discussed above.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

For instance, a parent/guardian and child could *both* go in-person to a third-party entity that assesses the existence of their relationship and their respective ages, and then generates a code indicating that parental consent has been given for the creation of the minor's account. That code could then be accepted by the first-party point of access as evidence that parental consent has been provided. Such a code could also permit the parent/guardian to quickly access the child's account if needed.

At present, many unknowns surround the implementation of such provisions. However, as policymakers grow increasingly interested in such requirements, it is likely that third-party firms will further invest in systems designed to validate the existence of parent-child relationships. For instance, a third-party age verification company might treat minors' accounts as derivatives of a primary account held primarily by an assigned "parent" figure, whose age is in turn verified according to one of the methods contemplated above.[31]

C. AUDIT PROCESSES

In addition to establishing age-verification requirements for certain online platforms, policymakers likely will wish to require federal or state regulatory bodies to monitor the entities involved for compliance. Among other relevant considerations, policymakers should specify the range of entities subject to compliance audits, the timetable for such audits, the procedures for undertaking such audits, and the level of compliance required for an entity to successfully complete an audit.

Additionally, any mandated audit procedures should consistently reflect the core institutional design choices discussed above—what kind of data is involved, and which entities are involved with its verification.

D. RESTRICTIONS ON DATA USE

Given the amount and sensitivity of data necessarily associated with an age-verification regime, policymakers should consider imposing strict controls on the use of such data by whichever entities retain it.

If a first-party verification model has been selected, platforms might be prohibited from selling any data associated with or derived from a user's demographic information not otherwise made available to the platform, and/or prohibited from using such data for behavioral-advertising purposes. For instance, a social media platform that uses images of a user's driver's license for age-verification purposes might be prohibited from cross-referencing the driver's license number against public databases, finding out that the user has a record of speeding tickets, and then marketing them defensive-driving courses on that basis.

Similar restrictions might be applied to third-party companies engaged in age verification, which might likewise be tempted to sell or otherwise monetize user data provided to them.

E. RESTRICTIONS ON DATA RETENTION

Depending on the verification mechanism selected, policymakers might impose additional controls on the extent to which "raw" user data might be retained by the entities involved.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

Any age verification regime involves the collection of data. On a first-party verification model, "raw" data would be collected by the entity operating the point of access, such as a social media platform. On a third-party verification model, "raw" data would be collected by the entity with which the first-party point of access contracts. On a zero-knowledge proof model, "raw" data would be collected by the trusted third party that produces the proof used to validate the user's encrypted data.

The question for policymakers is whether this "raw" data should be retained indefinitely, deleted after a period of time, or "forgotten" immediately after age is verified. Indefinite retention presents privacy concerns by raising the risks associated with a data breach. Conversely, a regime requiring the deletion of "raw" user data becomes essentially "un-auditable": a third-party regulator cannot review a subset of user accounts and the associated data to ascertain whether age verification did in fact occur consistent with the requirements of a given law or policy. The decision here must be made by policymakers according to the goals they seek to achieve.

The ideal compromise solution would probably bar the indefinite retention of user data, requiring deletion of age-verification information after one to three years or a sustained period of account inactivity. Legal provisions establishing audit mechanisms would need to be tailored to account for these realities.

F. PENALTIES FOR VIOLATIONS

For any age-verification regime to prove successful, policymakers must give it teeth—particularly where violations of its terms are concerned. Where entities mishandle user data, they should face stiff consequences.

Breaches incurring civil—or even criminal—penalties might involve outright noncompliance with a law or policy's terms, particularly where audit procedures are concerned, or involve negligence or malfeasance in the exposure of user data to hostile actors or other unauthorized personnel. Policymakers should also consider ensuring that the scope of such penalties is sufficiently broad so as to apply to any third-party entities involved in the collection of users' information.

Enforcement of age-verification laws might be carried out several ways. State attorneys general or departments of consumer protection could serve as primary government enforcers, with their efforts augmented by tools like private rights of action. Parents whose children are provided social media services apart from parental consent could be empowered to bring lawsuits—individually or through class actions—against providers who violate age-verification laws. While aggressive, measures like these would likely have a significant deterrent effect on bad actors.

Importantly, however, penalties should be sufficiently tailored to account for inevitable gaps in data sets. No age verification system will be 100% effective, and accountability measures built into any policy should reflect that.

Conclusion

In crafting any age-verification regime, policymakers must balance (1) user privacy, (2) corporate power, and (3) children's safety. As this memo has demonstrated, "age verification" is an expansive umbrella concept. Numerous different configurations of policies can be prescribed by legislators or

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

other policymakers, based on stakeholders' assessment of the relevant equities. Once core structural decisions have been made—the type of data used for verification, and the entity conducting the verification—that design can be expanded to address other relevant considerations, such as the preferred audit and penalty mechanisms.

The technology behind identity verification has already advanced to the point where quite sophisticated and robust models are now possible. Credit cards and digital government identity documents are now ubiquitous, and a number of identity verification technology companies are already seeing significant growth with private-sector applications.

However, it is still important to note that many of these approaches are novel and have not yet been tested in court. No matter how carefully they are undertaken, it is likely that any policymaking efforts in this space will be promptly met with First Amendment challenges spearheaded by well-funded corporations and aligned legal groups. Policymakers should prepare for this inevitability by minimizing legal exposure to constitutional attacks.

In light of these considerations, the age-verification model we recommend as the ideal approach would be (1) based on government identification documents, and (2) built around the use of a third-party service that would either use a traditional method of verification or a zero-knowledge proof to verify user age. Government ID numbers and other information can be checked against existing records databases—an important reliability safeguard—and need not be disclosed to first-party entities for effective age verification to occur. Nor do they present the pervasive privacy risks associated with the use of biometrics, which policymakers should avoid. This would also limit any risks associated with the monetization of user data by first-party entities engaged in advertising markets.

A next-best solution might require first-party entities to collect government ID information, but only when accompanied by an aggressive audit regimen on the parts of regulators and stiff penalties for mishandling of user data. While potentially less protective of user privacy, this approach would not require the collection of biometric information, would still likely be effective in deterring underage access, and would still allow enforcers to validate company compliance against existing public records databases.

The least-optimal approaches would be any approach involving the collection of biometric information—a maximally-invasive approach—or a reliance on inferential/indirect information, which (at the opposite extreme) would likely prove wildly unreliable as a proxy for user age.

In legislative findings and other messaging surrounding a law or policy structured according to this model, it is important to narrowly focus on regulation of covered platforms as regulation of certain instrumentalities of communication, rather than regulation of particular categories of user speech. This would increase the odds that such a law would survive the inevitable First Amendment scrutiny in light of *Ashcroft*.

Such a law or policy should also contain a robust audit protocol: data gathered by a third-party engaged in verification might be checked against public records databases for completeness and accuracy, and covered platforms could be held accountable for the third-party entities with which they contract. (Since it relies on audits for enforcement, this approach necessarily entails that data be retained by the entities involved, rather than deleted.) Penalties for systematic or willful noncompliance, on the parts of any entities involved, should also be included.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

Finally, a successful age-verification regime should be configured so as to permit parental access to, and control over, the accounts of minor children. As previously discussed, there are a number of ways in which such configuration might occur. As above, a third party-based solution—in which parents and children have their relationship and respective ages verified by an outside entity, who then validates that information for the first-party point of access—likely best balances procedural reliability with users' interest in data privacy.

Regardless of the precise strategy selected, it is highly likely that debates over social media's harms and benefits will continue well into the future—absent radical changes in how users relate to the platforms. Age verification, when properly designed and implemented, offers state policymakers an important tool for holding powerful technology companies accountable and keeping young people safe.[32]

---

[1] Twenge, J. M., Haidt, J., Joiner, T. E., & Campbell, W. K. "Underestimating digital media harm." *Nature Human Behaviour* 4 no. 4 (2020): 346–48.
[2] Pasquale, L., Zippo, P., Curley, C., O'Neill, B., & Mongiello, M. "Digital Age of Consent and Age Verification: Can They Protect Children?" *IEEE Software* 39 no.3 (2020): 51.
[3] Perez, S. "TikTok CEO Says Company Scans Public Videos to Determine Users' Ages." *TechCrunch*, March 23, 2023.
[4] Weatherbed, J. "Utah Governor Signs New Laws Requiring Parental Consent for Under-18s to Use Social Media," *The Verge*, March 24, 2023.
[5] Kelley, J., & Schwartz, A. "Age Verification Mandates Would Undermine Anonymity Online." *Electronic Frontier Foundation*, March 10, 2023.
[6] Chavez, K. "What Would Legally-Mandated Age Verification on the Internet Actually Look Like in Practice?" *NetChoice*, March 21, 2023.
[7] Drivers Privacy Protection Act of 1994, 18 U.S.C. § 2721 (authorizing bona fide business use of drivers' license databases for purposes of identity verification).
[8] 18 U.S.C. § 1028 (criminalizing the falsification of identification documents).
[9] Huddleston, J. "Would New Legislation Actually Make Kids Safer Online?" *Cato Institute*, April 6, 2023.
[10] Kelley, J., & Schwartz, A. "Age Verification Mandates Would Undermine Anonymity Online."
[11] Snow, J. "Why Age Verification Is So Difficult for Websites." *The Wall Street Journal*, February 27, 2022.
[12] Asaravala, A. "Why Online Age Checks Don't Work." *Wired*, October 10, 2002.
[13] Stokel-Walker, C. "Lying About Your Age? This AI Will See Right Through It."
[14] Op. Cit., Pasquale et al., "Digital Age of Consent and Age Verification: Can They Protect Children?."
[15] *Id.*
[16] Commission Nationale de l'Informatique et des Libertés, "Online Age Verification: Balancing Privacy and the Protection of Minors." *CNIL*, September 22, 2022.
[17] Stokel-Walker, C. "Lying About Your Age? This AI Will See Right Through It." *Coda*, March 28, 2023.
[18] Griswold, C. "Protecting Children from Social Media." *National Affairs,* Spring 2022.
[19] Kairaldeen, A.R.; Abdullah, N.F.; Abu-Samah, A.; Nordin, R. "Peer-to-Peer User Identity Verification Time Optimization in IoT Blockchain Network." *Sensors* 23 (2023)*: 2107.*
[20] Buchanan, William J. "Zero-knowledge Proof: Proving age with hash chains." Asecuritysite.com (2023).
[21] *Id.*
[22] Guerra, E., & Westlake, B. G. "Detecting child sexual abuse images: traits of child sexual exploitation hosting and displaying websites." *Child Abuse & Neglect* 122 (2021):105336.
[23] *Ashcroft v. ACLU*, 542 U.S. 656, 661 (2004).
[24] *Id.* at 662.
[25] *Id.* at 661. See also *Reno v. ACLU*, 521 U.S. 844, 874 (1997) ("In order to deny minors access to potentially harmful speech, the CDA effectively suppresses a large amount of speech that adults have a constitutional right to receive and to address to one another.").
[26] *Ashcroft*, 542 U.S. at 660 (citations omitted).

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.

[27] *Id.* at 662.
[28] *Id.* at 668.
[29] *Id.* at 667.
[30] *Id.* at 667 (emphasis added).
[31] Polonetsky, J. "Age Verification for Children: A Survey of Tools and Resources." *International Conference of Data Protection and Privacy Commissioners*, November 2009.
[32] While this report focuses on approaches that states can take, it is worth being aware that increasingly, policymakers at the federal level have been showing an interest in age-verification laws. To take just one notable and recent bipartisan example, in early 2023, Senators Brian Schatz, Tom Cotton, Chris Murphy, and Katie Britt introduced the *Protecting Kids on Social Media Act*, a bill which would—among other things—impose age verification requirements on social media platforms and create a federal pilot program for the provision of secure digital identification credentials. Protecting Kids on Social Media Act, S. 1291, 118th Cong. (2023); see also Making Age-Verification Technology Uniform, Robust, and Effective Act, S. 419, 118th Cong. (2023); Social Media Child Protection Act, H.R. 821, 118th Cong. (2023) (presenting alternative proposals for age verification). In essence, the pilot proposal in their bill would amount to a type of third-party verification (where the federal government would be act as the third-party entity involved). While this particular proposal has the advantage of potentially broad applicability, it is worth also noting that any universal digital identifier administered by the federal government raises long-term privacy questions.

© Copyright 2023 Institute for Family Studies and the Ethics and Public Policy Center. All rights reserved.