Yale Journal of Law & Technology
Volume 28, Symposium Issue

Online Age Gating: An Interdisciplinary Evaluation

Noah Apthorpe[*], **Brett Frischmann**[**], **Yan Shvartzshnaider**[***]

The recent surge in regulation seeking to establish age-based governance online is part of a decades-long attempt to establish online zoning. It is driven by active development of technologies to estimate or verify user age based on various characteristics of users, their credentials, or their activities. However, these developments have heightened prevailing concerns that online age gating technology will inevitably be abused and misused to cause a variety of privacy harms and rights infringements. This paper examines this ongoing debate by bridging technical and legal scholarship to explore the current state of online age-based governance. We discuss the current legal and policy landscape, the current status of online age gating technologies, and provide recommendations to guide legal and technological scholarship and practice. Our interdisciplinary assessment is particularly important and timely, given the recent flurry of state and federal laws that aim to implement age gating online and ongoing litigation challenging such laws.

[*] Department of Computer Science, Colgate University, Hamilton, N.Y.
[**] Charles Widger School of Law, Villanova University, Villanova, P.A.
[***] Department of Electrical Engineering and Computer Science, Lassonde School of Engineering, York University, Toronto, Canada.

Article Contents

Introduction

Governments have struggled to maintain and establish zoning online for decades. In some cases, the struggle has concerned exercising jurisdiction. Physical borders in the offline world generally work well (though not perfectly) in enabling state governance, but those borders do not necessarily map well onto the digitally networked environment.[1] In other cases, the struggle has involved attempts to utilize various technologies to establish boundaries based on the characteristics of users or their activities online so that legal or social norms might be recognized and enforced by states, communities, platforms, or others (e.g., employers, teachers, or parents). These zoning efforts have run into various obstacles, including other technologies used to attack, disrupt, or circumvent boundary creation, maintenance, or enforcement. Recent work on "governance seams," which explores how friction in the design of borders, boundaries, and interfaces can enable governance, emphasizes the role of competing sociotechnical influences on what is possible/feasible and what is desirable/legitimate.[2]

Consider *geoblocking*, for example. This practice involves imposing governance rules based on geolocation, often to establish and apply jurisdictional concepts to the online world. Initially, geoblocking online seemed impossible to do effectively, given technological limitations and the ease of circumvention. In the past 25 years, however, the situation has changed dramatically. Geolocation tracking is not only widespread and normalized, but there are many sophisticated means for reliably inferring someone's geolocation. Of course, the cat-and-mouse game (or arms race) persists, especially as virtual private networks (VPNs) and anonymity-focused overlay networks like Tor[3] become more accessible.

---

[1] David R. Johnson & David Post, *Law and Borders: The Rise of Law in Cyberspace*, 48 STAN. L. REV. 1367 (1996).

[2] Brett Frischmann & Paul Ohm, *Governance Seams*, 37 HARV. J.L. & TECH. 1117 (2023).

[3] Tor is a free, open-source anonymity network that uses layered encryption and routes Internet traffic through a series of relays, making it difficult for observers to trace users' identities, locations, or browsing activity.

Nonetheless, while not without tradeoffs, governance seams based on geolocation are feasible, widely employed, and broadly effective for users not expressly attempting circumvention.[4]

This paper focuses on another prominent example: *age gating*. Age gating involves imposing governance rules based on a person's age or age category. Offline, age gating is a common practice around the world.[5] Many jurisdictions apply governance rules based on the age of citizens. Age serves as a trigger for different governance rules in many different contexts for many different reasons. Familiar examples include age restrictions on access to alcohol, drugs, and pornography, as well as age criteria for voting, military service, driving licensure, contracting, and employment.

Age gating is a type of governance seam, in that some mechanism for assessing age generates friction at a sociotechnical border, boundary, or interface to enable some form of governance. The available mechanisms and associated frictions vary, as do the potential governance rules. It is important to clarify that age gating is not limited to binary allow/disallow rules. Governance rules can be quite nuanced and varied based on the context. For example, being below a particular age may trigger a parental consent requirement, set a privacy or security default, or prohibit access to services or content altogether. In simplest terms, age gating only means that age serves as the triggering condition for governance rules: if [age], then [governance rules].

Governments around the world have passed laws to implement age gating online.[6] In the United States, a flurry of age gating laws has been proposed or enacted by states and challenged in the courts. Many of these laws require content-based age gating, limiting access to online material based on

---

[4] Jack Goldsmith & Eugene Volokh, *The Relevance of Ross to Geolocation and the Dormant Commerce Clause*, 102 TEX. L. REV. ONLINE 30 (2023).

[5] *See*, *e.g.*, Laurence Steinberg and Grace Icenogle, *Using Developmental Science to Distinguish Adolescents and Adults Under the Law*, 1 Annual Review of Developmental Psychology 21, 26-27 (2019) (describing origins of age boundaries and various examples).

[6] *See*, *e.g.*, Australia, Online Safety Act 2021; United Kingdom, Online Safety Act 2023.

the content of the material. One prominent example is a 2023 Texas law (H.B. 1181) that requires age gating for "sexual material harmful to minors."[7] The Texas law is representative of similar laws in several other states.[8] Not surprisingly, these laws spark furious debate about free speech and the First Amendment.[9] While many online age gating laws are content-based, others are not. For example, New York State's 2024 Stop Addictive Feeds Exploitation (SAFE) for Kids Act prohibits the provision of social media feeds with addictive design patterns to minors regardless of content.[10] Still other laws incorporate both content-based and content-neutral provisions, such as California's 2023 Age-Appropriate Design Code Act (AADC)[11], the 2023 Federal Kids Online Safety Act (S. 1409)[12], and others, several of which are still under consideration, are currently under judicial review, or have been enjoined.

Online age gating has become such an active and politically charged issue and topic of extensive ongoing litigation that an interdisciplinary evaluation is needed. Legal and policy debates about age gating circle around three types of normative concerns: 1) the threat of *privacy* invasions and surveillance; 2) the overburdening of adult speech and blocking minors' access to legitimate speech; and 3) abuse of *power*, namely, claims that age gating technology will inevitably be abused and misused by authoritarian or misguided states or

---

[7] H.B. 1181, 88th Leg., Reg. Sess. (Tex. 2023).

[8] "H. B. 1181 is not the only law of its kind. At least 21 other States have imposed materially similar age-verification requirements to access sexual material that is harmful to minors online." *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461, slip at 3 (5th Cir. 2023).

[9] The Fifth Circuit upheld this Texas law in *Free Speech Coal., Inc. v. Paxton*, 95 F.4th 263, Free 95 F.4th 263 (5th Cir. 2023), and as this article went to press, the U.S. Supreme Court affirmed that decision in a 6-3 decision. *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461, slip at 3 (5th Cir. 2023). We had offered analysis of the Fifth Circuit majority and dissenting opinions and provided some recommendations for the Supreme Court. We keep most of that discussion, which remains relevant, and modestly update our analysis to reflect on lessons (not) learned.

[10] Stop Addictive Feeds Exploitation (SAFE) for Kids Act, N.Y. GEN. BUS. LAW §§ 1501-1508 (McKinney 2024).

[11] Age-Appropriate Design Code Act, CAL. CIV. CODE § 1798.99.28–40 (West 2023).

[12] Kids Online Safety Act, S. 1409, 118th Cong. (2023-2024).

parents/guardians.[13] All three types of concerns are often advanced without considering how they might be mitigated or eliminated by different technical implementations.[14] Despite the fact that the magnitude, scale, scope, and nature of privacy and speech concerns depend substantially on the technical specifics of age gating systems, these specifics are rarely examined. Too often, the debate among academics and policy makers, also reflected in amicus briefs, involves unsubstantiated claims about privacy, speech, and abuse of power that ignore technological details and consequently are divorced from empirical reality. This article aims to improve the quality of such debates.

We contend that there are two strong reasons to reevaluate the case for age gating online. First, progress in available means for online age determination might reduce many of the relevant burdens. Evaluating this requires rigorous frameworks for analyzing age determination technologies and various cost-benefit tradeoffs. Second, social demand for age gating online has changed over the past few decades. While early efforts focused predominantly on age gating adult content, and those efforts persist, society has increasingly recognized other legitimate reasons for age gating, such as imposing governance rules focused on product design (e.g., to set privacy-protective defaults or to ban socially harmful designs, such as infinite scrolling). Further, age gating need not only focus on minors; there may be compelling reasons to protect other vulnerable populations, such as the elderly, from

---

[13] Concerns about the abuse of power is the untouchable "third rail" of the age gating debate, often expressed as a slippery slope argument akin to "once you build a technology that can age gate for narrow governance rule X, it can and inevitably will be used for problematic governance rule Y and Z." Unfortunately, this argument too easily preempts evaluation and fails to engage both with (i) genuine issues of power, authority, and jurisdiction (who gets to decide? what use cases are authoritarian? who decides who is (mis)guided?) and (ii) technical designs that might mitigate the perceived risks.

[14] We reject tech solutionism and do not hold out technical design options as ready-made solutions to complex sociotechnical problems. Rather, we argue that the technical details matter to understand what is possible, frame comparative analysis, and evaluate tradeoffs that are often shaped by the technological systems deployed.

online manipulation and exploitation. Reasonably evaluating the normativity of age gating in light of changing social demands requires a better exploration of technical feasibility.

Accordingly, this paper bridges technical and legal scholarship to explore the current state of online governance based on user age. We review the existing literature and draw from our own experience as computer scientists and legal scholars. We first outline the history of online age gating, its connection to offline age gating, and compare content-based and content-neutral motivations for age-based governance (Section I). We discuss the current legal and policy landscape, focusing on representative high-profile laws and court cases from 1968 to 2025 (Section II). We present a categorization of online age gating technologies and an evaluation framework, including threat models and desired properties (Section III). We review the current state of online age gating technologies, including research efforts and deployed services, and reason about tradeoffs (Section IV). We synthesize our legal and technical analyses into recommendations for legislatures, regulatory agencies, and courts (Section V) and provide a brief conclusion.

I. Background

This section provides an overview of offline and online age gating, including age gating based on content, capabilities, and vulnerabilities.

A. Offline Age Gating

Age gating is widespread and works reasonably well in the offline context for two main reasons:

First, age can be self-authenticating offline with some reasonable degree of confidence to support the application of governance rules. As Lessig argued decades ago, physical appearance often provides a sufficient basis to deny most children access to pornography offline.[15] Beyond information transmitted through casual observation, offline social

[15] Lawrence Lessig, *The Law of the Horse: What Cyberlaw Might Teach*, 113 HARV. L. REV. 501, 503-04 (1999). Despite some false positives and false negatives, physical appearance is a pretty good signal of one's age.

interaction and communication also provide contextual information that supports informed age estimates. Asking someone their age and follow up questions can be reasonably effective in many offline scenarios. Of course, there are borderline cases where the reliability of this approach weakens. Nevertheless, physical appearance and social interactions provide useful information that reduce the costs of age authentication by narrowing the range of people for whom a more costly verification process is required. Further, the availability of such means channels actors towards certain interactions and potential threat models, a topic we revisit below.

Second, age authentication through government-issued identification (e.g., passports and driver's license) works well enough to support the application of governance rules. Many offline contexts, such as purchasing alcohol or cigarettes or renting a car, require the presentation of government-issued identification to verify an individual's attributes. So regular are these processes that the use of government issued identification for offline age gating is a reasonable baseline against which most age gating proposals should be compared. While some circumvention risks persist (e.g., fake IDs), the continued use and widespread acceptance of offline age gating with government-issued identification indicates that it performs well enough to support governance to a broad extent.

However, some raise serious objections, including that some forms of offline age gating disproportionately impact particular groups, and in some cases, groups that are already disadvantaged. For example, age gating that relies on government identification or other documentation, such as a bank statement, can impose substantial burdens that "fall disproportionately on the elderly, the poor, the undocumented, those living in rural areas, etc."[16]

Offline age gating can also disproportionately burden people who lose their documents in an emergency, such as a flood, and people for whom a trip to the DMV is incredibly

[16] Janet Zhang & Steven Bellovin, *Preventing Intimate Image Abuse via Privacy-Preserving Anonymous Credentials*, 26 SMU SCI. & TECH. L. REV. 149, 207 (2023).

difficult. It is critical to acknowledge and seriously account for these concerns when evaluating age gating designs in different contexts. However, these concerns should not preempt age gating proposals. There are tradeoffs, which are often context specific, and there are different means for addressing disproportionate burdens, which also may be context specific.[17]

B. Online Age Gating

Despite the mundanity of offline age gating, online age gating has not been widely adopted. Early efforts (circa 1990s-2010s) failed, seeding the widespread understanding that effective online age gating would be difficult for a variety of sociotechnical reasons. Technological limitations and the ease of circumvention played important roles, as did other important considerations, including concerns about privacy and burdening adult speech.[18] To make matters complicated, these factors are interdependent, and all seem to depend on the design of the Internet. The Internet architecture provides incredible affordances to users that expand the scale, scope, and reach of digital communications and corresponding economic and social interactions across the globe, but it also makes reliable authentication of various attributes (geolocation, age, identity, etc.) and corresponding application of governance rules more difficult.[19]

We make four general points here about how the Internet architecture shapes the technological feasibility and tradeoffs of online age gating:

First, the Internet protocol (IP) and other protocols used to interact with Internet services (TCP, UDP, HTTP, etc.), do not require devices to provide information about the identity or age of users. Rather, they associate devices, and by proxy, users, with pseudonymous identifiers (MAC addresses, IP

---

[17] *Id.* at 208.

[18] Jonathon W. Penney, *Chilling Effects: Online Surveillance and Wikipedia Use*, 31 BERKELEY TECH. L.J. 117 (2016); Elizabeth Stoycheff et al., *Privacy and the Panopticon: Online Mass Surveillance's Deterrence and Chilling Effects*, 21 NEW MEDIA & SOCIETY 602 (2019).

[19] For further discussion of how Internet architecture makes reliable authentication difficult and the intelligence-disaffordances of end-to-end design, see Lessig, *supra* note 15; Brett Frischmann & Blake Reid, *Network Neutrality as Governance Seam*, TECHREG CHRON., May 2024, at 1.

addresses, port numbers) that can only be linked with real world identities through separate processes, e.g., requiring a website account sign-in that associates a session with a specific user. Thus, age gating and other forms of identity-based governance must necessarily be "bolted-on" to a fundamentally pseudonymous system. It could have been otherwise if core Internet protocols had required participating users to provide intelligence necessary for age authentication along with a trustworthy process for verifying such information. However, doing so would have undermined core design principles and corresponding user affordances, social commitments and values, such as democratized speech opportunities and pseudonymous access.[20] This historical design contingency matters for an interdisciplinary evaluation of online age gating because it shapes the relevant tradeoffs, including privacy baselines, speech opportunities, and potential chilling effects.

Second, the Internet is not a centralized entity. It is a network of independent autonomous systems that control their own membership and access rules. This distributed design means that no single network or collection of networks has sufficient authority to authenticate the identities or ages of all Internet users. Similarly, individual websites and online services are typically responsible for their own user authentication processes. Attempts to incorporate new age gates into this existing system must therefore be decentralized by design with due consideration of jurisdictional boundaries and the limited power of participating entities.

Third, the Internet pushes intelligence and functionality to programmable end-hosts, ranging from consumer devices (laptops, smartphones, etc.) to servers in massive data centers. Users have broad freedom to customize the behavior of these devices, either by writing custom code or by acquiring and running existing software tools and hardware peripherals. As the past few decades have demonstrated, software running on increasingly powerful user devices lowers the costs of many different types of authentication circumvention; digital "disguises" are more varied, accessible, and effective for a

---

[20] *Id.*

wider range of users.[21] This matters for an interdisciplinary evaluation of online age gating because it complicates the analysis of threat models, and it drives an arms-race because authentication systems must also operate on end-host software.

Fourth, Internet users are distributed physically with connected individuals and organizations located around the world. For offline age gating,[22] actors are in the same geographic location, social setting, and legal jurisdiction. For online age gating, this need not be, and often is not, the case. This fact introduces various complications. As Lessig emphasized, age is not self-authenticating online.[23] Further, social interactions and communications may be a less reliable source of age authenticating information when people are not present in the same place with shared cultural contexts. Finally, different communities might set different age thresholds, prescribe different rules, have different normative priorities, and tolerate different trade-offs.

While age may be more easily authenticated offline than online, the relative paucity of online age gating was neither preordained nor fundamental to the concept of an online environment. To our knowledge, there has never been a serious per se objection to age gating online. Nor was there serious objection to the idea that the state could have a legitimate interest in age gating online; for example, courts and commentators generally recognized the compelling state interest in restricting minors' access to pornography both offline and online. Instead, there have been technical/empirical debates about feasibility and reliability, and there have been contentious social/policy debates about the desirability and legitimacy of online age gating in light of various tradeoffs and potential spillovers, such as unintended external effects on privacy or speech.

We now present three high-level motivations for online age gating that encompass many of the good faith rationales provided in this space.

---

[21] *See infra* section 4.3.1 (discussing threat of user circumvention).
[22] Lessig, *supra* note 15.
[23] *Id.*

1. Content-based Age Gating

Many conversations about online age gating focus on *content*, usually adult content and attempts to prevent minors from accessing pornography. Calls for content-based age gating have also included appeals to block LGBTQ content,[24] educational content,[25] video games, and violent or graphic content,[26] among others. These appeals often elicit vehement counterarguments from those concerned about the infringement of First Amendment rights; authoritarian surveillance, censorship, and/or repression; public/private power imbalances; and other ills that could arise from limiting access to information.[27]

Unfortunately, this furor has left two other major age gating motivations without sufficient attention and consideration from the legal and technical community: age gating based on *capabilities* and age gating based on *vulnerabilities*. As we describe below, we take capability-based and vulnerability-based age gating to be content-agnostic and focused on factors such as legally defined capacities and the susceptibility of minors and older adults to exploitation, manipulation, and social engineering techniques that appear in all varieties of online spaces.[28] We appreciate that some may use the language of capabilities and vulnerabilities to motivate age gating that is

---

[24] TUMBLR STAFF, TUMBLR (July 18, 2013), https://staff.tumblr.com/post/55906556378/all-weve-heard-from-a-bunch-of-you-who-are [https://perma.cc/V9HD-EAE3].

[25] ACLU MASSACHUSETTS, *Amid Growing Calls for Censorship in Schools and Libraries, Massachusetts Advocates and Educators Support Bill to Protect Students' Right to Learn*, ACLU MASS. (Jan. 10, 2024), https://www.aclum.org/en/press-releases/book-bans-censorship-massachusetts-schools-libraries [https://perma.cc/PKJ7-2YQQ].

[26] Hayden Smith & William Cipolli, *The Instagram/Facebook Ban on Graphic Self-Harm Imagery: A Sentiment Analysis and Topic Modeling Approach*, 14.1 POLICY & INTERNET 170 (2022).

[27] Matthew P. Hooker, *Censorship, Free Speech & Facebook: Applying the First Amendment to Social Media Platforms via the Public Function Exception*, 15.1 WASH. J.L. TECH. & ARTS 36 (2019).

[28] In some cases, manipulation harms and content-based harms are separable and independent, and in other cases, they are inseparable and interdependent. *See* Brett Frischmann and Peter Ormerod, *Regulation of manipulative design is not preempted by CDA 230 or the First Amendment*, EMORY L. J. (forthcoming 2026).

content-based, for example, when someone argues that minors are vulnerable to violent video games or pornography Ultimately, in such situations, we would categorize the motivation as content-based rather than capability- or vulnerability-based.

2. Capability-based Age Gating

Most communities agree that minors lack the capability (autonomy) to perform various activities of modern life, including signing contracts, giving consent, waiving legal rights, voting, and serving in the military. This may also be true for some adults; for example, the elderly suffering from dementia. Heightened restrictions of these activities triggered by age have nothing to do with content or censorship. Instead, the restrictions generally concern a person's competencies, development, maturity, or other attributes that are, or at least are perceived by the community to be, required for the relevant activity. Age often serves as a proxy for competency. Of course, some minors may be more capable (competent, mature) than many adults, but most are not.

Consider, for example, contracts. Laws restricting minors' ability to enter binding contracts or to waive legal rights are content-neutral. One rationale for such age gating is that minors' relevant capacities develop during childhood and adolescence, and minors generally lack the competence to enter binding contractual commitments.[29] Minors' brains perform differently than adults when making decisions, reasoning, and solving problems.[30] Such differences can lead adolescents and children to be more impulsive and less deliberative and can impact information processing and

---

[29] As described by one comment on the common law origins of the infancy doctrine, "The reason for allowing the minor the privilege of voiding his contracts was the protection of the minor. It was thought that the minor was immature in both mind and experience; therefore, he should be protected from his own bad judgment as well as from adults who would take advantage of him. There can be little dispute with the reasoning behind this purpose." James L. Sivils, Jr., Comment, Contracts-Capacity of the Older Minor, 30 U. KAN. CITY L. REV. 230, 230 (1962).

[30] *See, e.g.*, Steinberg and Icenogle, *Using Developmental Science to Distinguish Adolescents and Adults Under the Law*, 1 Annual Review of Developmental Psychology 21 (2019) (reviewing developmental science literature).

interpretation of cues, signals, and terms.[31] In the offline context, age gating contracts is widespread and not controversial.[32]

Yet every day, minors enter millions of digital contracts (such as end-user license agreements) and waive their rights repeatedly when using digital networked technologies. The fact that many of these contracts and waivers are not technically enforceable or at least can be "disaffirmed" under the infancy doctrine does not seem to matter. As Frischmann and Vardi explore,[33] digital contracting yokes minors and adults into networked relations with countless strangers through multi-sided markets that are the backbone of the surveillance economy. Frischmann and Selinger worry this will train or habituate minors to a scripted click-upon-cue behavior, setting expectations and teaching norms to accept and follow.[34]

This state of affairs is nonsensical. The development of simple, effective, and privacy preserving online age gating motivated by *capabilities* could align online spaces with existing capability-based laws that are widely accepted, content neutral, and relatively well-enforced offline. To continue the example, age gating could be required for digital contracting. This would protect all stakeholders, preventing those incapable of entering

---

[31] *Id.*; *See also* RESTATEMENT OF THE LAW, CHILDREN & LAW TENT. DRAFT 5-20.20.

[32] In most states, minors' freedom to contract is curtailed by the infancy doctrine, which allows minors to disaffirm contracts and render them unenforceable. *See* RESTATEMENT (SECOND) OF CONTRACTS § 14 (1979); RESTATEMENT OF THE LAW, CHILDREN AND THE LAW § 20.20 (Am. L. Inst., Tentative Draft No. 5, Mar. 2023); Robert G. Edge, *Voidability of Minors' Contracts: A Feudal Doctrine in a Modern Economy*, 1 U. GEORGIA L. REV. 205, 207-08 (1967).

[33] Brett Frischmann & Moshe Y. Vardi, *Better Digital Contracts with Prosocial Friction-in-Design*, 65 JURIMETRICS 1, 30–33 (2025). *See also* Michelle A. Sargent, *Misplaced Misrepresentations: Why Misrepresentation-Of-Age Statutes Must Be Reinterpreted as They Apply to Children's Online Contracts Authors*, 112 MICH. L. REV. 301, 308–09 (2013) ("In a world where contracts are formed effortlessly and where there is essentially no opportunity to bargain or negotiate, children making agreements online are more vulnerable than ever and require the protection of the infancy doctrine.").

[34] BRETT FRISCHMANN & EVAN SELINGER, RE-ENGINEERING HUMANITY 62–67 (2018).

into a legal agreement from doing so and preventing other parties from later learning that the contract is unenforceable. Of course, the feasibility, effectiveness, and tradeoffs of such a system would depend heavily on the details of the technical implementation and the governance practices involved, but the underlying capability-based motivation demands serious investigation from the technical and legal communities.

3. Vulnerability-based Age Gating

Minors and older adults are more vulnerable to exploitation, manipulation, and social engineering in various offline contexts. Few object to laws and social norms designed to protect them. Of course, there are exceptions, examples and contexts within which minors and older adults demonstrate more resilience and less vulnerability (e.g., teenagers capable of spotting a scammer/poser/adult in various online fora) or where stereotypes rather than evidence about these groups may drive attempts to age gate. Nevertheless, offline vulnerability-based age gating is generally accepted and reasonably effective. Again, contract law provides a decent example. A second common rationale for age gating contracts is a concern about exploitation: essentially, that adults will take advantage of minors.[35]

It seems reasonable to hypothesize that minors and older adults may be more vulnerable to data-driven manipulation, social engineering, addictive interface design (e.g., infinite scrolling), and "dark" and "not-so-dark-but-still-harmful" patterns online.[36] There is growing evidence that social media and smartphone use are contributing to a teen mental health epidemic[37] and many broader debates about relationships

---

[35] *See supra* note 29 and accompanying text.

[36] Christoph Bosch et al., *Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns*, 4 PROC. ON PRIV. ENHANCING TECHS. 237 (2016).

[37] *See, e.g.,* U.S. OFF. OF THE SURGEON GEN., SOCIALMEDIA AND YOUTHMENTAL HEALTH: THE U.S. SURGEON GENERAL'S ADVISORY 8 (2023) https://www.hhs.gov/sites/default/files/sg-youth-mental-health-social-mediaadvisory.pdf [https://perma.cc/74V4-H4K9]; Betul Keles, Niall McCrae & Annmarie Grealish, *A Systematic Review: The Influence of Social Media on Depression, Anxiety and Psychological Distress in Adolescents*, 25.1 INT'L J. ADOLESCENCE & YOUTH 79 (2020); Matthew Lawrence, Brett Frischmann, & Avi Sholkoff, *Tort Liability for Failure to*

between technology use and social well-being.[38] The scale and scope of positive and negative effects of various digital networked technologies and the variability of vulnerability across different populations is an active area of research.

Nonetheless, there is ample social demand for age gating that could protect vulnerable individuals online. Such protections could be content neutral, as they could permit, block, or modify design and user interface features regardless of the content these features present (and even if the content remains constant in response to design changes).[39] The case for vulnerability-based age gating online has grown stronger as society has begun to recognize compelling normative ends beyond policing access to content. As with capability-based age gating, vulnerability-based age gating demands serious consideration from the legal and technical community.

4. General Observations

Age gating for content-based, capability-based, or vulnerability-based rationales creates a governance seam that not only triggers specific governance rules but also allocates, enables, and even protects the exercise of governance authority by certain entrusted parties, such as parents or other guardians. Laws requiring parental consent, for example, may rely on one or more of the three rationales. These laws depend upon age gating, the effectiveness of which is especially important in situations in which honoring parental rights is constitutionally required.[40] Again, the shift from offline to online age gating presents a series of challenges and opportunities that demand interdisciplinary investigation and attention to social and technical details.

We acknowledge that relying on age as the triggering condition for governance rules motivated by capability- and vulnerability-based reasons is a second-best approach. Perhaps

---

Age Gate: A Promising Regulatory Response to Digital Public Health Hazards, J. TORT L. 1 (2025). https://doi.org/10.1515/jtl-2025-0018.

[38] Frischmann & Selinger, *supra* note 34.

[39] *See* Brett Frischmann and Peter Ormerod, *Regulation of manipulative design, is not preempted by CDA 230 or the First Amendment*, Emory L.J. (forthcoming 2026).

[40] *See* Pierce v. Society of Sisters, 268 U.S. 510 (1925).

the best approach would be to use capability and vulnerability as the triggering condition. That is, rather than verifying age and relying on age as a proxy, one might verify whether a person is capable or vulnerable and then apply the appropriate governance rules. Consider contracting, for example. Some 14-year-olds might be fully capable of contracting, and some 22-year-olds might not be. Using 18 as a fixed boundary age for contracting will therefore lead to both Type I and Type II errors and corresponding error costs.[41] Instead of relying on age, directly testing capability might reduce error costs, but it would greatly increase administrative costs and be prone to different types of errors. Again, we emphasize that these tradeoffs are contingent, contextual, and subject to change and reevaluation as technology advances. Age serves as a reasonably reliable proxy for capacity and vulnerability in many jurisdictions and social contexts, although there is considerable variety in what specific age serves as the cutoff.

II. Current Law and Policy

Here, we provide a high-level summary of legal and policy precedents that often frame debates about age gating online. Unsurprisingly, these focus mostly on age gating adult content and First Amendment challenges. As emphasized in the previous section, this provides a rather distorted view of the legal and policy considerations. Nonetheless, it is the common ground from which we start. We then turn our attention to the approaches and language used in recently proposed and enacted laws to give a sense of how lawyers and policymakers are conceptualizing age gating online. Finally, we identify a knowledge gap where a deeper understanding of age gating technologies and a framework for comparatively evaluating system designs would be useful.

In the offline context, states regulate minors' access to adult content, and courts generally have upheld such regulations.[42]

---

[41] *See, e.g.*, Jonathan B. Baker, *Taking the Error Out of "Error Cost" Analysis: What's Wrong with Antitrust's Right*, 80 ANTITRUST L.J. 1, 4 (2015)(explaining Type I and Type II errors and error costs)

[42] *See Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461, slip at 3 (5th Cir. 2023); Ginsberg v. New York, 390 U.S. 629 (1968) (applying rational basis review to uphold N.Y. statute that criminalized sale of adult content to minors).

Of course, this is just one of many, mostly uncontroversial, examples of age gating offline. Other examples include gambling, smoking, vaping, guns, voting, military service, and voting.

In the online context, states have tried to regulate minors' access to adult content, but those efforts have often failed; courts have held the laws unconstitutional.[43] Perhaps the most notable attempt to age gate adult content online is the Communications Decency Act of 1996 (CDA).[44] The law criminalized the knowing transmission of "obscene or indecent" content and the knowing sending or displaying of "patently offensive"' content to a person under 18 years of age.[45] The law provided two affirmative defenses, one of which covered those who take "good faith, reasonable, effective, and appropriate actions" to restrict access by minors and one of which covered those who use certain forms of age gating, such as a "verified credit card, … adult access code, or adult personal identification number."[46]

In *Reno v. ACLU*, the U.S. Supreme Court struck down these provisions of the CDA for violating the First Amendment.[47] Applying strict scrutiny because the law directly targeted the content of speech, the Court first recognized that the state has "a compelling interest in protecting the physical and psychological well-being of minors" and then evaluated whether the law was narrowly tailored to that end or instead was overbroad, vague, and overly restrictive of constitutionally protected speech.[48] The court examined many reasons for striking down the prohibitions. Most importantly, the statute included vague and ambiguous language that likely would chill

---

[43] *See* Reno v. ACLU, 521 U.S. 844 (1997) (holding portions of Communications Decency Act unconstitutional); Ashcroft v. ACLU, 542 U.S. 656 (2004) (holding Child Online Protection Act unconstitutional). *But see* United States v. Am. Libr. Ass'n, 539 U.S. 194 (2003) (upholding Children's Internet Protection Act).

[44] Communications Decency Act of 1996, Pub. L. No. 104-104, 110 Stat. 133 (1996), *invalidated in part by* Reno v. ACLU, 521 U.S. 844 (1997).

[45] *Id.*

[46] *Id.* § 223(d)(e).

[47] *Reno*, 521 U.S. 844.

[48] *Id.* at 869.

protected speech,[49] particularly given criminal penalties, and the prohibitions could have been more narrowly tailored in various respects (e.g., by limiting its scope to commercial speech).[50]

Among other arguments, the government argued that the age gating defense narrowed the statute's burden on speech.[51] The Court acknowledged that age gating by use of "a verified credit card or adult identification" is not only "technologically available but actually is used by commercial providers of sexually explicit material.[52] These providers, therefore, would be protected by the defense."[53] The Court nonetheless rejected the government's argument, reasoning:

> Under the findings of the District Court, however, it is not economically feasible for most noncommercial speakers to employ such verification. Accordingly, this defense would not significantly narrow the statute's burden on noncommercial speech. Even with respect to the commercial pornographers that would be protected by the defense, the Government failed to adduce any evidence that these verification techniques actually preclude minors from posing as adults. … the District Court correctly refused to rely on unproven future technology to save the statute.[54]

The two relevant factors weighed by the Court were both heavily contingent on the then-existing state of technology. First, economic feasibility for non-commercial speakers depended (mostly) on the cost of available age verification technology.[55] Second, the threat of minors posing as adults depended on the availability, cost, and effectiveness of

---

[49] *Id.* at 872.
[50] *Id.* at 846. Technically, the controlling opinion says it does not reach the overbreadth inquiry, although it discusses the issue briefly.
[51] *Id.*
[52] *Id.* at 881.
[53] *Id.*
[54] *Id.* at 881-82.
[55] *Id.* at 880.

circumvention technology, given the available age verification technology.[56] These factors are thus dynamic and relative, rather than absolute. While the fact findings of the district court determined the outcome in that case, the relevant facts (and thus factors to be weighed) have changed substantially over the past few decades, as we discuss below. The question whether age gating technologies are now "proven," economically feasible, and capable of "saving" state-mandated governance rules is open.

Taking into account the various failings of the CDA identified by the Court in *Reno*, Congress drafted the Child Online Protection Act ("COPA") to shore up many of the definitional and scope issues in the CDA and reinstate age gating of adult content.[57] In a fractured opinion, the Supreme Court struck down COPA, on the view that filtering and blocking software was a less restrictive alternative to COPA's age-based zoning.[58] Much of the Court's analysis turned on factual claims about the existing state of different technologies.[59] We leave further discussion aside, but note that COPA had a slightly broader affirmative defense for "any provider who in good faith, has restricted access by minors to material that is harmful to minors— (A) by requiring use of a credit card, debit account, adult access code, or adult personal identification number; (B) by accepting a digital certificate that verifies age; or (C) by any other reasonable measures that are feasible under available technology."[60] Again, the Supreme Court did not talk in absolute terms or set forth per se rules; instead, it compared the relative burdens of different then-available technologies.[61]

In 2024, a Fifth Circuit Court of Appeals decision upheld a Texas law similar to the section of the CDA struck down by *Reno*.[62] The Texas law regulated "commercial entities that

---

[56] *Id.* at 882.
[57] Child Online Protection Act of 1998 (COPA), 47 U.S.C. § 231, *invalidated by* Ashcroft v. ACLU, 542 U.S. 656 (2004).
[58] *Ashcroft v. ACLU*, 542 U.S. 656, 657 (2004).
[59] *Id.* at 668.
[60] *Id.* at 662 (quoting Child Online Protection Act of 1998 (COPA)), 47 U.S.C. § 231, *invalidated by* Ashcroft v. ACLU, 542 U.S. 656 (2004).
[61] *Id.* at 668.
[62] *Free Speech Coal., Inc. v. Paxton*, 95 F.4th 263 (5th Cir. 2024).

knowingly and intentionally publish or distribute sexual material on an Internet website, including a social media platform, more than one-third of which is sexual material harmful to minors."[63] The law required these entities to use age verification to limit the material to adults and display health warnings on the landing page of the website and advertisements.[64]

In analyzing the constitutionally relevant burden of the age verification requirement, the Fifth Circuit did not apply strict scrutiny and follow the path set by the *Reno* decision; instead, it applied rational-basis review, relying on *Ginsberg*, which involved a New York state law that age gated sex-related content in the offline context.[65]

In *Ginsberg*, the Supreme Court decided the regulated material was obscene for minors and thus not protected by the First Amendment, but for adults, the regulated material was not obscene and thus constitutionally protected speech.[66] The Court approved this two-tier approach, applied rational basis review of the state law, and upheld it.[67] Though not explicitly stated by the Court (or even germane to the issue it directly confronted), the Court's application of rational basis review in *Ginsberg*[68] was supported by basic technological facts; age gating to separate minors and adults worked reasonably well and was normal in the offline, in-person context. The government could directly regulate content deemed obscene for minors with a rather limited, but still real, burden on adult access to protected speech. Those burdens were nothing more than the conventional costs of age authentication in many familiar offline settings, such as liquor stores, bars, and other retail establishments, whether implemented at the door to the premises or at the checkout counter. This zoning worked in

---

[63] *Id.* at 267 (quoting Publication or Distribution of Sexual Material Harmful to Minors on an Internet Website; Providing a Civil Penalty, 2023 Tex. Sess. Law Serv. Ch. 676 (codified at Tex. Civ. Prac. & Rem. Code Ann. § 129B.001 et seq.)).
[64] *Id.* at 266.
[65] *Id.* at 278.
[66] *Id.* at 269.
[67] *Ginsberg v. New York*, 390 U.S. 629 (1968).
[68] *Id.* at 639.

large part because authenticating age in a brick-and-mortar store was relatively cheap, effective, and common.

As we have stressed, those basic technological facts change dramatically when moving from offline to online age gating, and the facts continue to change with technological advancement. A detailed technical evaluation of such changes is thus critically important.

In contrast with the reasoning in *Reno*, and more in line with the reasoning and factual premises of *Ginsberg*, the Fifth Circuit majority reasoned that "the age-verification requirements do not impose any sort of 'categorically different' burden on adults" than the burden adults face offline.[69] The law requires regulated entities to use "reasonable age verification methods" including use of "government-issued identification," "digital verification," or other "commercially reasonable method[s]."[70] The court reasoned further:

> That allows for at least three concrete means of age-verification: (1) government ID, (2) facial appearance, or (3) some other available information used to infer the user's age. At least one of those options will have no more impact on privacy than will in-person age verification a la *Ginsberg*.[71]

Thus, the Fifth Circuit majority concluded that Texas could age gate online effectively without imposing greater burdens on adults than they would face in offline age gating contexts.[72]

Judge Higginbotham dissented on this point, contending that Texas's attempt at age gating online necessarily imposed additional burdens on adult speech that should trigger strict scrutiny.[73] Specifically, Judge Higginbotham stated that "the New York statute at issue in *Ginsberg* did not burden the free speech interests of adults, but H.B. 1181 [the Texas law] does [because it] requires that adults comply with the age

---

[69] *Free Speech Coal.*, 95 F.4th at 271.
[70] *Id.* at 300 (Higginbotham, P., concurring in part).
[71] *Id.* at 271.
[72] *Id.*
[73] *Id.* at 299 (Higginbotham, P., concurring in part).

verification procedure and view the required health disclosures before accessing protected speech."[74]

Yet, as noted above, the New York statute in *Ginsberg did* impose a burden, if limited, on adults. Adults had to authenticate their age to legally purchase covered content. Perhaps Judge Higginbotham believed this burden did not burden free speech interests, but it is hard to see how that can possibly be the case, as the extra effort, time, and disclosure of private information alone presumably would, on the margins, affect some adults' ability and willingness to purchase the content. A better explanation is the one we suggested above, which is that the age authentication burden present in *Ginsberg* was conventional and normalized. This does not mean that it does not exist or that the magnitude of the burden is zero. Instead, it means that the friction associated with the governance seam is simply an expected and accepted part of the background social environment. After all, there is no First Amendment right to a frictionless speech environment, which is why in many cases, courts evaluate state interventions (e.g., new laws) in terms of the additional burdens, given the background environment.

As our article went to press, the Supreme Court issued an opinion affirming the Fifth Circuit decision.[75] The Supreme Court applied intermediate scrutiny and found the Texas law constitutional.[76] This precedent will shape age gating jurisprudence for the foreseeable future. Thus, after reviewing the technological landscape in Sections III and IV, we revisit the Fifth Circuit analysis and examine the Supreme Court's decision in more detail.

In recent years, another push for age gating laws involves *age-appropriate design codes*, such as California's Age-Appropriate Design Code (AADC)[77] and Maryland's Kids Code (also known as the Maryland Age Appropriate Design Code)[78], and *laws implementing age gates for addictive feeds*,

---

[74] *Id.* at 293 (Higginbotham, P., concurring in part).
[75] *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461 (2025).
[76] *Id.*
[77] California Age-Appropriate Design Code Act, CAL. CIV. CODE §§ 1798.99.28-40 (West 2025).
[78] Maryland Age-Appropriate Design Code Act, MD. CODE ANN., COM. LAW §§ 14-4801-4813 (West 2024).

such as California's Protecting Our Kids from Social Media Addiction Act[79] and New York's Stop Addictive Feeds Exploitation (SAFE) for Kids Act.[80] These approaches shift away from content-based to capability- and vulnerability-based age gating. Given concerns over how design engineers behavior, these laws aim to shift design patterns from antisocial to prosocial behaviors. Litigation challenging these and other child safety laws is ongoing and thus far has been largely successful in stalling enforcement.

A key takeaway from past legal challenges to state efforts to age gate online is that the state of technology shapes the relevant tradeoffs, including the relative burdens on adults who must interact with age gates. Indeed, the Supreme Court has expressly recognized the technological contingency of its decisions regarding age gating. Nevertheless, the intersection of legal scholarship and technical evaluation of age gating systems remains underexplored. Legal and technical experts have tended to talk past each other on this issue, making unsubstantiated, conclusory claims about online age gating, including about its technological feasibility or impossibility and its supposedly inevitable substantial burdens on privacy and speech.

We seek to address this gap by providing a more rigorous framing and examination of the sociotechnical facts and tradeoffs surrounding online age gating. We are not committed to a particular view about age-gating technology and are generally skeptical of technological solutionism. We also believe minors and the elderly may be particularly vulnerable to data-driven techno-social engineering and that safeguards like privacy-protective defaults and restrictions on addictive or other harmful designs are potentially defensible as content-neutral age gating regulations, depending very much, however, on the details concerning the underlying technology.

---

[79] Protecting Our Kids from Social Media Addiction Act, CAL. HEALTH & SAFETY CODE §§ 27000-27007 (West 2025).
[80] GEN. BUS. LAW §§ 1501-1508.

III. Age Gating Technology Overview

The design and implementation of online age gating technologies is an active area of research and development. Age gating is governance of the format if [age], then [governance rules], and most recent technological advances have focused on the if [age] component, i.e., how to determine whether a user meets an age threshold. This is a challenging problem given the nature of the Internet (Section I) and the potential for burdens and tradeoffs (Section II). We correspondingly focus our survey of age gating technologies here and in Section IV on methods of online age determination (if [age]).

This section provides a high-level categorization of age gating technologies followed by a set of considerations for evaluating and comparing age determination systems, including threat models and sociotechnical factors, that can provide relevant insight into the legal and technical implications of these technologies. These considerations are applicable to all age gating technologies regardless of whether they arise from content-based, capability-based, or vulnerability-based motivations. Section IV then reviews a wide variety of age gating technologies ranging from well-understood and widely deployed to more theoretical approaches undergoing active research and development.

A. Categorization

Online age gating technologies can be divided into three broad categories based on their method of age determination:

- Self-attestation gates ask users to confirm or deny whether they meet an age threshold. This can involve asking users for their exact age, asking users whether they are over (or under) a specific age, asking a user to provide their birth date for comparison to a target date, or similar queries. Self-attestation age gates are widespread and often implemented via a simple checkbox, date entry, or set of yes/no buttons.[81]

[81] Tajveer Singh Dhesi & Noah Apthorpe, *Measuring the Prevalence and Variety of Online Age Gates*, 2025 IEEE WORKSHOP ON TECHNOLOGY AND CONSUMER PROTECTION (2025).

- Age estimation gates use information about users to estimate, or *infer*, their age. This inference can be performed using a variety of possible methods, ranging from simple logic to complex machine learning models. The information underlying the inference can similarly be of a wide variety of types and provenances. Age estimation is an active area of research that is being increasingly incorporated into deployed age gating systems.[82] The user experience can vary widely, and the corresponding burdens and tradeoffs are also highly variable depending on the design of the system.
- Age verification gates use information about users in conjunction with another trusted source or authoritative verification service to attempt to *confirm* their age. Unlike self-attestation gates or age estimation gates, age verification gates attempt to obtain some independent, trusted verification that the determined age of the user (or whether the user meets an age threshold) is correct. Age verification gates are also a topic of active research attempting to improve functionality and mitigate various burdens.

Websites or online services may deploy age gating technologies from multiple categories. For example, a service that typically relies on age estimation might flag some users for a more thorough age verification process.

B. Considerations for Evaluating and Comparing Age Determination Systems

As age gating technologies receive greater attention, more widespread deployment, and heightened research efforts, it will be increasingly important to have structured approaches to evaluate technical designs and tradeoffs. As discussed in

---

[82] YOTI LTD., YOTI FACIAL AGE ESTIMATION WHITE PAPER (2024), 1, 2 (2024), https://cdn.aws.yoti.com/wp-content/uploads/2026/01/Yoti-Age-Estimation-White-Paper-July-2025-PUBLIC-v1.pdf [https://perma.cc/5A9Z-XJS3]; James Beser, *Extending Our Built-In Protections to More Teens on YouTube*, YOUTUBE BLOG (2025), https://blog.youtube/news-and-events/extending-our-built-in-protections-to-more-teens-on-youtube [https://perma.cc/L5S4-WTZQ]; OpenAI, *Building Toward Age Prediction* (2025), https://openai.com/index/building-towards-age-prediction [https://perma.cc/4T42-S4KP].

Section II, the U.S. Supreme Court has not set forth per se rules about age gating; instead, it has compared relative burdens of different available technologies. We therefore present a series of evaluation considerations that can be used by a variety of stakeholders to compare age determination (if [age]) technologies. We do not claim completeness; other considerations may arise in particular contexts or with particular system designs. Rather, we intend to provide a common framework for productive discussions and debates about age determination technologies. We focus on the if [age] step and do not engage with the full range of normative considerations appropriate for evaluating whether and how to apply governance rules at the then [governance] step. Such considerations would venture beyond technical system design.

C. Threat Models

Evaluating and comparing age gating technologies requires consideration of the possible threats to which they may be subjected. The cybersecurity literature often enumerates a list of "threat models," each defining the goals and formal capabilities of a particular actor seeking to undermine a system. While we do not strive for the formality of a technical security paper, we consider the following threats essential for consideration when evaluating online age gating technologies.

For brevity, we focus on active threats in which one or more human (or organizational) actors may be unreliable or behave arbitrarily. We do not consider threats unrelated to core age determination functionality, programming bugs, or numerous non-active threats facing most technological systems.

Each relevant threat could reduce the effectiveness of age gating technologies or place undue burdens on users, including chilled speech, de-anonymization, profiling, or surveillance.

1. Circumvention

Users are often strongly incentivized to circumvent online age gates, as the governance rules they enforce may restrict access to resources or actions that are desirable to those who do not meet required age thresholds.[83]

---

[83] Fatmaelzahraa Eltaher et al., *The Digital Loophole: Evaluating the Effectiveness of Child Age Verification Methods on Social Media*,

Consider a user who enters an incorrect birthday on an online questionnaire, a user who records and replays network traffic corresponding to an older user's interaction with a website, a user who uploads a photo of themselves that has been modified using an generative artificial intelligence tool to make them look older, a user who employs a virtual private network (VPN) to pretend they are location in a different jurisdiction, or a user who "borrows" their parent's driver's license or credit card to submit the associated number into a website form. Each of these, and many other possible circumvention attempts, may or may not be successful depending on the technological details of the age gating system.

Users are not the only parties that may seek to circumvent age gates. Online services that are compelled to deploy age gates by regulation may subtly or overtly attempt to influence the age determination process to allow users to bypass governance rules. For example, a service might tweak the parameters of a machine learning model such that it more likely produces estimates above an age threshold. Another service might apply design patterns or nudging text that push users toward actions that allow them to bypass an age gate and away from actions that would provide actual knowledge that they do not meet the age threshold.

All online age gating technologies must therefore be evaluated with minimal assumptions about the motivations, methods, and technical sophistication of potentially circumventing users and Internet services. All else being equal, an age gating technology that is less prone to circumvention is preferable; however, designing to prevent circumvention often introduces other tradeoffs, including several of the threats below.

2. Scope Creep

The threat of scope creep is a concern about actors who use age determination technologies for purposes other than the particular if [age] determination that is legitimately part of an age gate. Concern about scope creep is often raised as an

---

INTERNATIONAL CONFERENCE ON INFORMATION SYSTEMS SECURITY AND PRIVACY (2025).

objection to online age gating because extending the use of the age gating tool may be difficult to track and may lack justification.[84]

The governance seams created by online age gates are quite narrow by definition. Any action taken or enabled by an age gate other than a strict if [age], then [governance rules] pattern is an example of scope creep. Unfortunately, online age gates are very susceptible to scope creep, as they introduce a new interaction point between a user and an online service that may involve new data transfers. More sophisticated types of age estimation and age verification gates are especially prone to scope creep, since they involve more complex interactions and data that are more relevant for other uses, such as images of users' faces, credit card numbers, or government IDs.

A particularly concerning example of scope creep is the combination of age gating with *identity verification*, a similar but ultimately unrelated governance seam. Identity verification seeks to obtain and confirm various personally identifying details of a user, such as their name, birthdate, address, ID numbers, etc., and then take actions based on the fact of the user's identity. The most important distinction between identity verification and age gating is that identity verification is necessarily de-anonymizing, while age gating is not.

If a hypothetical age gating technology were to learn only a single bit of information about a user—whether the user meets an age threshold or not—this technology would be able to apply all forms of age-based governance. In contrast, an identity verification technology must learn many bits of information about a user—enough to uniquely identify them among all people within some jurisdiction—to effectively apply identity-based governance. However, since data used for identity verification also likely provides enough information for age determination, there is a natural inclination toward the merging of these two technologies. Unfortunately, many commercially available age gating services are repurposed

[84] Scope creep is the idea that a tool designed and justified for one purpose gradually ends up being used for other purposes, which may or may not be appreciated, much less justified. *See* Frischmann and Selinger, *supra* note 34, ch.3.

identity verification services that learn and retain more about users than is strictly necessary for age gating.[85]

Many online services face incentives to combine age gating with other actions, such as building profiles of users for advertising or surveillance. For example, a government service used to verify a user's age might also seek to record a link between the user and the age-gated content or actions they attempt to access. A website might use a photo provided for age estimation to infer a user's demographics and tailor their experience on the website to their inferred race or social class. Such potential is rightly highlighted by critics of online age gating and sometimes used as a slippery slope argument in calls to eliminate online age gating entirely.[86] While we share these concerns, we note that this threat is not *fundamental* to online age gating. Rather, it is a result of external incentives, including online monetization strategies, Internet design practices, and the current data privacy regulatory regime.[87]

Actual and potential scope creep must be considered in all meaningful evaluations of age gating technologies when determining the burdens they impose on users. While not all forms of scope creep are negative (e.g., an age gate that also serves as a friendly welcome to a website), we are skeptical, by

---

[85] *E.g.*, iDenfy, *Best Age Verification Software Providers of 2025*, https://www.idenfy.com/blog/best-age-verification-software-providers-of-2024/ (last visited Jul. 1, 2025) [https://perma.cc/92L9-JSWN]; Secure Identity, LLC, *Simplifying Age Verification For Businesses and Consumers*, https://identity.clearme.com/post/simplifying-age-verification-for-businesses-and-consumers (last visited Jul. 1, 2025) [https://perma.cc/V7AP-DF8E].

[86] Electronic Frontier Foundation, *Age Verification Harms Users of All Ages*, https://www.eff.org/files/2024/10/21/age_verification_-_policy.pdf (last visited Jul. 1, 2025) [https://perma.cc/5337-HV2E]; Martin Sas & Jan Tobias Mühlberg, *Trustworthy Age Assurance?*, THE GREENS CLUSTER: SOCIAL & ECONOMY, LOCATION: THE EUROPEAN PARLIAMENT (2024).

[87] *See* SHOSHANA ZUBOFF, THE AGE OF SURVEILLANCE CAPITALISM: THE FIGHT FOR A HUMAN FUTURE AT THE NEW FRONTIER OF POWER 5 (2019); JULIE E. COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM (2019); BRETT FRISCHMANN & EVAN SELINGER, REENGINEERING HUMANITY 63 (Cambridge Univ. Press 2018); Brett Frischmann & Susan Benesch, *Friction-in-Design Regulation as 21st Century Time, Place and Manner Restriction*, 25 YALE J.L. & TECH. 376, 387 (2023).

default, when age gating is combined with other processes, especially processes that cause or enable de-anonymization or profiling. The technical and legal community should reflect this skepticism when considering the expansion of online age gating and explore whether design and regulation can produce online age gating methods that are resistant to scope creep and/or are verifiably limited to if [age], then [governance rules].

3. Data Misuse

As noted above, age gating technologies need to learn and communicate only a single bit of information about a user in order to apply if [age], then [governance] rules. In practice, correctly learning that single bit of information typically requires the use of more rich data sources. As we describe in Section IV, current and proposed age gating technologies draw from a wide variety of data to determine user age, including images, web browsing patterns, government IDs, credit cards, and more.

Services that deploy or support age gating technologies may decide to collect more information than needed, store that information for longer than needed, share or sell that information with other entities, or use that information for a purpose other than to make an age determination. Some of these actions overlap with scope creep, but they are all essential design considerations when evaluating the tradeoffs posed by age determination technologies. Data that is unnecessarily collected or stored poses a risk to users due to the increased possibility of data breach or later secondary use, whether anticipated by the designers of the technology or not.[88]

Data collection and storage can be tricky to identify, since modern web services often incorporate user data into training sets for artificial intelligence models, which retain the data in their trained parameters even after the original data may be deleted.[89] Similarly, user profiling services often use machine learning to make new inferences about users that are based on,

---

[88] *See* Daniel Solove, *A Taxonomy of Privacy,* 154 U. PA. L. REV. 477 (2006).
[89] Congzheng Song et al., *Machine Learning Models that Remember Too Much*, PROCEEDINGS OF THE 2017 ACM SIGSAC CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY (2017).

but not exactly present in, the original user data.[90] Finally, data handling practices are often quite opaque to users, courts, and regulators, as data can be stored, sold, and shared entirely in and between back-end databases and business-to-business transactions that cannot be audited with front-end access or Internet monitoring tools.[91]

Privacy scholars have raised serious concerns about data misuse as a negative externality of online age gating.[92] As privacy researchers, we concur that data misuse is a serious threat posed by age determination technologies. Forms of data misuse directly constitute certain privacy harms and increase the risk of others.[93] If users know, expect, or believe that age gating will result in excessive data collection, long-term data storage, data sale or exfiltration, or a greater potential for surveillance, profiling, or other forms of secondary use, they may also be less inclined to engage with the gated experience, potentially raising First Amendment concerns.[94]

However, as with scope creep, we also note that data misuse is not *fundamental* to online age gating and is an artifact of the current technological and regulatory landscape. Age determination technologies should be evaluated based on their adherence to data handling best practices endorsed by privacy advocates and their vulnerability to external pressures—such as profit motives or surveillance demands—that could compromise those practices over time. Furthermore, the technical and legal community should explore whether new designs and regulations can produce online age determination

---

[90] *E.g.*, Amazon.com, Inc., *Amazon's AI-powered 'Interests' Feature Automatically Finds New Products that Match Your Passions and Hobbies*, https://www.aboutamazon.com/news/retail/artificial-intelligence-amazon-features-interest (last visited Jul. 1, 2025) [https://perma.cc/NX5H-VWHU].

[91] Matthew Crain, *The Limits of Transparency: Data Brokers and Commodification*, 20 NEW MEDIA & SOCIETY 88 (2018); Ashley Kuempel, *The Invisible Middlemen: A Critique and Call for Reform of the Data Broker Industry*, 36 NW. J. INT'L L. & BUS. 207 (2016).

[92] John T. Cross, *Age Verification in the 21st Century: Swiping Away Your Privacy*, 23 J. MARSHALL J. COMPUTER & INFO. L. 363 (Winter 2005); Sarah Scheffler, *Age Verification Systems Will Be a Personal Identifiable Information Nightmare*, 67.7 COMMUNICATIONS OF THE ACM (2024).

[93] *See* Daniel Solove, *supra* note 86.

[94] Jonathon W. Penney, *Chilling Effects: Online Surveillance and Wikipedia Use*, 31 BERKELEY TECH. L.J. 117 (2016).

methods resistant to data misuse that can protect vulnerable populations without sacrificing user privacy.

D. Design Features and Sociotechnical Factors

Rigorously evaluating age gating technologies requires consideration of various *positive* features in addition to the potential risks described above. We identify six features below that should be considered when conducting a rigorous technical evaluation of online age gating technologies. The first three of these features (correctness, unforgeability, and non-transferability) concern the ability of a technology to make effective age determinations that are resilient to various forms of circumvention. The next two of these features (unlinkability and data minimization) concern the reduction of privacy risks due to scope creep and data misuse that could have substantial consequences for individuals and society, including chilling effects on online participation, erosion of trust in online services, and the amplification of systemic inequalities when vulnerable individuals are exposed to privacy harms. The final feature (usability) concerns user experiences when interacting with age gates and the minimization of interaction costs.

All else being equal, age determination technologies that perform better across these features are preferable. However, these considerations are nuanced and non-binary, and tradeoffs exist when these features conflict due to real world constraints. We do not prescribe how tradeoffs between these features should be weighed, as their relative importance will depend on the specifics of the situation. It is the responsibility of legislatures, courts, and other democratic processes to make such decisions.

Although *transparency* is not listed here, it is an essential prerequisite, in that it is impossible to evaluate an age determination technology on the other features if actual technical details (e.g., source code, training data, and protocol specifications) are not accessible for audit.[95] Many commercial age determination services do not provide sufficient transparency into their source code or data for review. In such

---

[95] *See* Edwin A. Farley & Christian R. Lansang, *AI Auditing: First Steps Towards the Effective Regulation of Artificial Intelligence Systems*, 38 HARV. J.L. & TECH. DIG. 1 (2025) (many parallels between audits of AI systems and audits of age gating systems).

cases, those considering the potential burdens of age gating mediated by these services should assume the worst.

Importantly, strong performance on the below features does not de facto justify the use of an age determination technology. Rather, this list is intended as a continuation of our structured framework for comparing technologies and reasoning about technological tradeoffs in order to inform political and judicial deliberations about online age gating more generally.

1. Correctness

Given accurate input information, age determination technologies should be able to accurately determine whether a user's age meets a relevant threshold. The importance of this property is self-evident, as it prevents age-related governance rules from being incorrectly applied to users of the wrong ages. However, this property is also non-trivial, as age is not self-authenticating online.

Importantly, the correctness of a particular technology can vary considerably across users and situations. This could depend on users' actual ages (e.g., more accurate for children and seniors than for teenagers), users' technology use patterns (e.g., more accurate for heavy Internet users), or other user demographics (e.g., more accurate for users with ID information in more than one government database). Age estimation technologies are particularly susceptible to unpredictable correctness errors, although age verification technologies may also encounter correctness challenges depending on design and implementation choices.

There are various ways to measure the correctness of age determination technologies. A simple metric, accuracy, measures the ratio of correct determinations to all attempted determinations. While intuitive, accuracy does not always provide enough information to understand correctness tradeoffs, motivating the use of additional metrics, including precision, recall, true positive rate (TPR), false positive rate (FPR), true negative rate (TNR), false negative rate (FNR),

and mean absolute error (MAE).[96] It can be difficult to assess the correctness of age determination technologies analytically, so many rely on empirical measurements instead. Unfortunately, empirical measurements are highly sensitive to user populations and are vulnerable to various forms of bias, including intentional manipulation meant to artificially inflate the apparent performance of the technology. For example, the same technology, designed to determine whether users are older than 18, might perform fantastically if tested with users who are all either 10 or 30 years old but terribly if tested with users who are all either 17 or 19 years old. As this example implies, transparency is essential for reasoning about the relevance and reliability of correctness metrics.

2. Unforgeability

It should be impossible (or at least quantifiably difficult) for users to forge whatever input information is used by an age determination technology to determine their age.[97] Unlike correctness, which refers to a technology's ability to accurately determine user age from correct inputs, unforgeability refers to a technology's resistance to deceptive or otherwise incorrect input information. Unforgeability protects a technology against circumvention attempts from users or other entities.

3. Non-transferability

Users should be unable (or at least strongly disincentivized) to share credible input information used by age determination technologies among themselves, and technologies should be able to detect and reject transferred information that does not

---

[96] *Precision* is the proportion of users approved as meeting an age threshold who actually meet the threshold. *Recall* and the *true positive rate (TPR)* are the proportion of users who truly meet the age threshold that are correctly approved. The *false positive rate (FPR)* is the proportion of users who do not meet the age threshold but are incorrectly approved. The *true negative rate (TNR)* is the proportion of users who do not meet the age requirement and are correctly denied. The *false negative rate (FNR)* is the proportion of users who meet the age requirement but are incorrectly denied. The *mean absolute error (MAE)* is the average difference (e.g., in years) between users' actual ages and the ages determined by the system.

[97] Jonathan Katz & Moti Yung, *Unforgeable Encryption and Chosen Ciphertext Secure Modes of Operation*. INTERNATIONAL WORKSHOP ON FAST SOFTWARE ENCRYPTION (2000).

pertain to a specific user.[98] Like unforgeability, this property protects against circumvention and prevents the technology from inappropriately applying governance rules based on faulty or deceptive input. For example, a technology with strong non-transferability would be able to detect if a minor submits a government ID "borrowed" from a parent and reject the credential, while a technology with strong unforgeability would be able to detect and reject a "fake" ID with user-generated or user-modified details. While these properties are similar, they often require different algorithms and system designs.

4. Unlinkability

According to a strict formulation of unlinkability, age determination technologies should prevent the linking of any information collected about a user for a single instance of age determination to any other past, present, or future information about the user.[99] This means that the technology should make it impossible to tell whether information used for age determination is related in any way to information about the user from other contexts or from other interactions with the age gate. Unlinkability defends against scope creep and data misuse, as it prevents age gating from being used to build or contribute to profiles of users[100] and prevents collusion by organizations to identify that different credentials, actions, or content views belong to the same user.

Unlinkability is challenging in practice, since the types of data naturally relevant for age determination (e.g., birth dates,

---

[98] Gerrit Bleumer, *Biometric Yet Privacy Protecting Person Authentication*, INTERNATIONAL WORKSHOP ON INFORMATION HIDING (1998); Russell Impagliazzo & Sara Miner More, *Anonymous Credentials with Biometrically-Enforced Non-Transferability*. PROC. OF THE 2003 ACM WORKSHOP ON PRIVACY IN THE ELECTRONIC SOCIETY (2003); Sebastian Pape, *A Survey on Non-Transferable Anonymous Credentials*, IFIP SUMMER SCH. ON FUTURE OF IDENTITY IN THE INFO. SOC'Y 107 (2008).

[99] Sandra Steinbrecher & Stefan Köpsell, *Modeling Unlinkability,* PRIVACY ENHANCING TECHS.; 3D INT'L WORKSHOP 32 (2003); Mayla Brusó et al., *Linking Unlinkability,* 7TH INT'L SYMPOSIUM ON TRUSTWORTHY GLOBAL COMPUTING 129 (2012).

[100] Anna Lysyanskaya et al., *Pseudonym Systems*, SELECTED AREAS IN CRYPTOGRAPHY: 6TH ANN. INT'L WORKSHOP 184 (2000).

government IDs, face images, etc.) are inherently highly identifying and are readily linked to other datasets of user behavior online. Even the network or browser metadata involved in an interaction with an age gating service (e.g., IP addresses, packet timings, user agent strings, etc.) could be enough to fingerprint users and link age gating to other online activities.[101] Eliminating the use of any linkable data from age gating technologies is unlikely, so cryptographic protocols that prevent data linking by specific untrusted entities involved in an age determination system would need to be combined with regulatory and auditing frameworks to prevent data misuse.

5. Data Minimization

According to a strict formulation of data minimization, age determination technologies should collect only information needed to determine whether a user's age meets the threshold for an associated governance rule, delete the data as soon as that determination has been made, not share or sell the data, and not use the data for any other purpose.[102] Data minimization also implies that user age should be determined to the least specificity needed by the associated governance rule (e.g., "older than 13") and that collecting exact user ages should be avoided if unnecessary. Data minimization defends against scope creep and data misuse, as it places guardrails on access to user data.

Data minimization is not as strict as anonymization, because it still allows identifying data to be collected. Age determination technologies that do provide complete anonymity to users would be preferable to those with less comprehensive guarantees, but acceptable practices may exist at different levels of data minimization.

Data minimization is not a new concept. More or less strict versions have been incorporated into consumer data privacy laws, including the European General Data Protection

---

[101] Laperdrix et al., *Browser Fingerprinting: A Survey*, 14.2 ACM TRANS. ON THE WEB 8 (2020).

[102] *See* European Data Protection Supervisor, *Glossary*, https://www.edps.europa.eu/data-protection/data-protection-glossary_en (last visited Jul. 1, 2025) [https://perma.cc/GPL3-LEAA].

Regulation (GDPR)[103] and the California Consumer Privacy Act (CCPA).[104]

Comparing the data minimization practices of different age determination technologies should consider what components of a technological system have access to user data. For example, it might be more acceptable if user data is collected and processed only by a local device (e.g., a smartphone) than if the data is transmitted to a cloud server or stored in a remote database. Similarly, data minimization should extend to all entities that handle user data as part of the age determination process, including business partners, government verification services, and the providers of incorporated software libraries.

Data minimization is challenging in practice, as online age gates, like all technological systems, involve a complex web of interlinking technologies and managing entities. It may be difficult to trust one or more of these technologies or entities for a variety of sociotechnical reasons, including misaligned incentives or the risk of data breach. Intentional design can reduce the scope of who and what needs to be trusted, but there is unlikely to be a purely technical solution. Instead, technical designs that prioritize data minimization and reduce the trust surface will need to be combined with industry norms and regulatory oversight.

6. Usability Considerations: Interaction Costs and User Friendliness

We next present two related usability considerations that have implications for the design of age determination systems: interaction costs and user friendliness. There are likely others, as the usability field is quite extensive. We touch on these mainly to highlight the types of useability issues that should be considered when comparing the implications of potential age gating technologies.

Most age determination technologies introduce a new interaction step into the user interfaces of online services.[105] This interaction step can occur at various points, such as upon

---

[103] 2016 O.J. (L 679).

[104] California Consumer Privacy Act of 2018 (CCPA), CAL. CIV. CODE §§ 1798.100-1798.199.100 (West 2023).

[105] Tajveer Singh Dhesi & Noah Apthorpe, *supra* note 81.

site arrival, on account creation, or during checkout. This interaction step may only occur once for each user (e.g., if confirmation that the user meets a threshold age is associated with their account) or every time a user engages with the service. This intrusion of an age gate into the normal user experience poses an interaction cost in both time and effort that must be considered when comparing age determination technologies.

The higher the interaction cost imposed by the age gate, the more likely it is that the user will choose to leave the online service, only engage with elements of the service that are not age gated, or attempt circumvention of the age gate.[106] If the interaction cost is high enough that users are reasonably deterred from all or part of the online service, this may raise First Amendment concerns about chilling effects. As discussed in Section II, the time and effort required to engage with an age gate has been viewed as important in various cases regarding online and offline age gating. We suggest that empirical work on the actual, rather than speculative, chilling effects of different age determination technologies could test different levels and types of interaction costs.

User friendliness[107] is a related criterion regarding the time or effort required for user to engage with age-based governance compared either against a baseline of no age-based governance or against the time or effort typically expected for other online governance seams. User friendliness also incorporates whether the technology is accessible, easy to use, and effectively communicates what information it collects, how that information is used, and its actual performance on all the properties in this section. Insisting on user friendliness reduces interaction costs, to a degree, and makes it less likely that users to whom age-based governance rules do not apply experience chilling effects.

Unfortunately, there seems to be no widespread agreement what constitutes "expected" or "appropriate" friction for an age-based governance seam online. Additionally, different users experience different subjective burdens from different

---

[106] *See generally* STEVE KRUG, DON'T MAKE ME THINK, REVISITED: A COMMON SENSE APPROACH TO WEB USABILITY (2013).

[107] *See generally* JAKOB NIELSEN, USABILITY ENGINEERING (1993).

user interfaces based on personal histories, expectations, demographics, accessibility needs, and other individual differences. All else being equal, interfaces that require fewer clicks, fewer data submissions, fewer interactions with multiple services, more robust accessibility features, less time to obtain an age determination, etc., may be preferable, but there are many practical tradeoffs that must be considered when judging user friendliness and how user friendliness should be weighed against the other properties above.

IV. Current Age Gating Technology

This section examines the current state of online age gating technology across the three categories presented in the previous section and from the perspective of the threat models and sociotechnical factors from our evaluation framework. We draw from the current research literature as well as descriptions of deployed or proposed technologies by industry groups and government entities. We do not claim comprehensiveness—our goal is rather to provide a primer on the current state and challenges of online age gating technologies.

A. Self-Attestation

Self-attestation age gates are widely deployed[108] and have become even more common in recent years due to increased regulatory pressure. In general, users interact with self-attestation gates by indicating their age or birth date, often combined with an agreement to a terms of service document. This approach has been widely used to prevent online services from having "actual knowledge"[109] of users under 13 years old so that they are not subject to U.S. COPPA restrictions on data collection and handling. Self-attestation is also widely used to age gate adult content from users under 18.[110]

---

[108] Martin Sas & Jan Tobias Mühlberg, *Trustworthy Age Assurance?* THE GREENS CLUSTER: SOC. & ECON., LOCATION: EUR. PARL. (2024).
[109] 16 C.F.R. § 312.2 (2025) (defining "Website or online service directed to children").
[110] Christine Marsden, *Age-Verification Laws in the Era of Digital Privacy,* 10 NAT'L SEC. L.J. 210 (2023).

Singh Dhesi and Apthorpe found that nearly all site arrival age gates on e-cigarette websites use self-attestation to determine whether users are allowed to access the site.[111] This included several varieties of self-attestation gates, including those with age-labeled buttons (e.g., "I am over 18"), those with birth date entry fields, those with a button to indicate a non-specific age over the threshold (e.g., "I am of legal age"), and those that do not present an actual "gate" but have a statement saying that the user indicates they are of legal age by continuing to engage with the website.[112]

Of course, self-attestation age gates leave much to be desired. While nominally correct as long as the user provides an accurate age, they are completely vulnerable to user circumvention through credential forging, as users can state whatever age or birth date they wish with no verification.[113] While misrepresenting one's age is often a violation of terms of service, there is usually no practical consequence. Additionally, most self-attestation age gates usually provide no guarantees about unlinkability or data minimization, meaning that that users' ages and birth dates could easily be incorporated into other profiles of their online behavior.[114]

Self-attestation age gates are also susceptible to user interface dark patterns[115] and "not-so-dark-but-still-harmful" patterns that incentivize user circumvention. For example, buttons that allow users to indicate that they meet a threshold age are often brightly colored, prominently placed, and emphasized as the preferred option. On the contrary, buttons that allow a user to indicate that they do not meet an age threshold are often gray, small, or otherwise de-emphasized.

While self-attestation age gates may stop users who arrived at gated websites or services accidentally (e.g., by clicking an ambiguous search result), self-attestation age gating is effectively just security theater, since the barrier for user

---

[111] Tajveer Singh Dhesi & Noah Apthorpe, *supra* note 81.
[112] *Id.* at 2.
[113] Christine Marsden, *supra* note 110.
[114] *See* Sarah Scheffler, *supra* note 92.
[115] Colin M. Gray et al., *The Dark (Patterns) Side of UX Design.* PROCEEDINGS OF THE 2018 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS (2018).

circumvention is so low.[116] The continued reliance on self-attestation age gates should serve as a reminder that the more sophisticated forms of age gating discussed below remain far from widespread. As some of the following technologies begin to see broader deployment, further research is needed to monitor longitudinal and cross-industry adoption and corresponding practical impacts.

B. Age Estimation

Age estimation has received substantial research focus in recent years, especially with interest in increasingly sophisticated machine learning techniques. Most of these research projects focus on age estimation alone and not on the application of age estimation to age gating (although many point out that age gating is a natural use case of their estimation method). We consider various age estimation techniques by dividing them into subcategories based the type of information they use to make age estimates.[117]

1. Biometrics

Biometric age estimation techniques attempt to infer user age based on biometric markers, such as facial features, voice patterns, and gait.

Age estimation based on facial features is a common age estimation approach in academic research and commercially available products. Most recent academic approaches use convolutional neural networks and related modern computer vision techniques to analyze images of users' faces and categorize them into age groups.[118] The reported accuracies of

---

[116] On security theater, *see* Bruce Schneier, BEYOND FEAR: THINKING SENSIBLY ABOUT SECURITY IN AN UNCERTAIN WORLD (Springer, Heidelberg, 2003); Bruce Schneier, *The psychology of security*, in INTERNATIONAL CONFERENCE ON CRYPTOLOGY IN AFRICA, pp. 50-79 (Springer Berlin Heidelberg, 2008).

[117] Azfar Adib et al., *Improvising Age Verification Technologies in Canada: Technical, Regulatory and Social Dynamics*, 2023 IEEE INTERNATIONAL HUMANITARIAN TECHNOLOGY CONFERENCE (2023).

[118] Abhinav Anand et al., *Age Estimation Based on Face Images and Pre-Trained Convolutional Neural Networks*, 2017 IEEE SYMPOSIUM SERIES ON COMPUTATIONAL INTELLIGENCE (SSCI) (2017); Yuntao Shou et al., CZL-CIAE: CLIP-Driven Zero-Shot Learning for Correcting Inverse Age Estimation (Feb. 27, 2014), https://ar5iv.labs.arxiv.org/html/2312.01758

these estimators vary considerably due to widely varying datasets (in number, quality, diversity, repetition, and orientation of face images), problem constructions (estimating exact ages or age groups, age alone or age plus other attributes), and methods (core machine learning algorithm, neural network architecture, training and testing hyperparameters, etc.).

Commercial products claim especially high accuracy. For example, the customer verification company Yoti reports that their "True Positive Rate (TPR) for 13-17 year olds correctly estimated as under 21 is 99.3%" and "TPR for 6-12 year olds correctly estimated as under 13 is 99.5%".[119] These metrics were determined by the National Institute of Standards and Technology (NIST) when Yoti submitted its model to NIST's Face Analysis Technology Evaluation (FATE) Age Estimation & Verification (AEV) track.[120]

The NIST FATE-AEV program tests machine learning models for age estimation based on face images by "execut[ing] the software on an open-ended set of different image databases that are appropriate to various use cases. Primary interest is in subjects making deliberate cooperative presentations to a

---

[https://perma.cc/V9MT-ZB2P]; Raphael Angulu et al., *Age Estimation via Face Images: A Survey*, EURASIP J. ON IMAGE & VIDEO PROCESSING 1 (2018); Zhenzhen Hu et al., *Facial Age Estimation with Age Difference*, 26 IEEE TRANS. ON IMAGE PROCESSING 3087 (2016); Hongyu Pan et al., *Mean-Variance Loss for Deep Age Estimation From a Face*, 2018 IEEE/CVF on COMPUTER VISION & PATTERN RECOGNITION 5285 (2018); Guodong Guo et al., *Human Age Estimation Using Bio-Inspired Features*, 2009 IEEE CONFERENCE ON COMPUTER VISION & PATTERN RECOGNITION 112 (2009); Wei Shen et al., *Deep Regression Forests for Age Estimation*, 2018 IEEE/CVF CONFERENCE ON COMPUTER VISION & PATTERN RECOGNITION 2304 (2018); Prachi Punyani et al., *Neural Networks for Facial Age Estimation: A Survey on Recent Advances*, 53 ARTIFICIAL INTELLIGENCE REVIEW 3299 (2020); Olatunbosun Agbo-Ajala & Serestina Viriri, *A Lightweight Convolutional Neural Network for Real and Apparent Age Estimation in Unconstrained Face Images*, 8 IEEE ACCESS 162800 (2020).

[119] YOTI LTD., *supra* note 82.

[120] *Id.* at 3. NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (NIST), *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*, https://pages.nist.gov/frvt/html/frvt_age_estimation.html (last visited Apr. 16, 2025) [https://perma.cc/7XEF-7TY2].

camera."[121] FATE-AEV is a follow-up to a 2014 NIST program with a similar goal.[122] Comparing the 2014 and 2024[123] NIST reports shows considerable improvement in age estimation performance over time, a trend also apparent from the academic literature referenced above.

However, all age estimation approaches based on face images have limitations that question their practicality for age gating. Not least among these are substantially reduced accuracy when users are in close proximity to relevant age thresholds. For example, it is effectively trivial to estimate that a 5-year-old is under 13. However, it is very challenging to estimate whether someone who is 17.9 (or 18.1) is over or under 18. The NIST FATE-AEV results reflect this difficulty (Figure 1[124]). Even the best models still estimated that 50% of

---

[121] Kayee Hanaoka et al., *Face Analysis Technology Evaluation Ongoing Age Estimation and Verification (AEV) Application Programming Interface (API), Version 2.0,* NAT'L INST.STANDARDS & TECH. (NIST) (Apr. 11, 2025), https://pages.nist.gov/frvt/api/FATE_AgeEstimation_API_v2.pdf [https://perma.cc/869F-T6UG].

[122] Mei Ngan & Patrick Grother, *Face Recognition Vendor Test (FRVT) - Performance of Automated Age Estimation Algorithms*, NAT'L INST.STANDARDS & TECH. (NIST) (Mar. 20, 2014), https://nvlpubs.nist.gov/nistpubs/ir/2014/NIST.IR.7995.pdf [https://perma.cc/2R47-Y8LS].

[123] Kayee Hanaoka et al., *Face Analysis Technology Evaluation Age Estimation & Verification,* NAT'L INST.STANDARDS & TECH. (NIST) (May 2024), https://pages.nist.gov/frvt/reports/aev/fate_aev_report.pdf [https://perma.cc/MCH7-5E8E].

[124] *Id.*

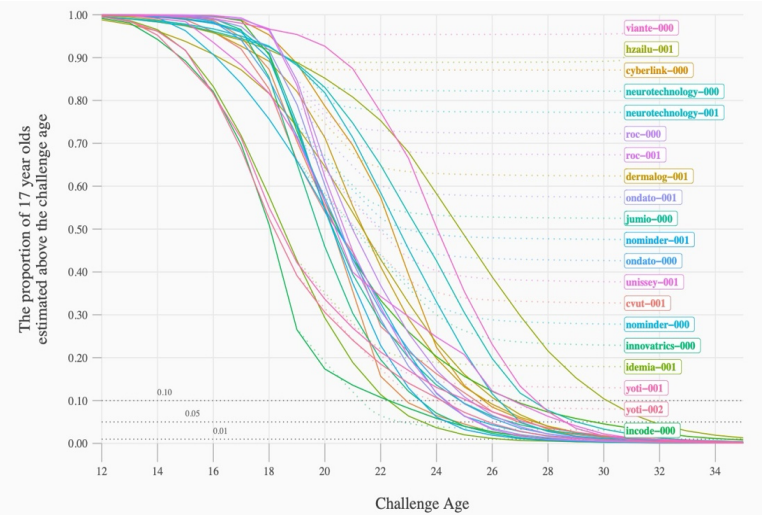

Figure 1. Figure from 2024 NIST Face Analysis Technology Evaluation (FATE) Age Estimation & Verification report: "The proportion of 17 year olds whose age is estimated as at or above the given challenge age."

17-year-olds were above 18, and the majority of models estimated that 50% of 17-year-olds were over 20.[125]

Low accuracy around age thresholds matters, because incorrectness in either direction could mean either that a legal adult is subjected to governance limiting their rights or a minor evades governance intended to keep them safe. As long as governance rules have sharp age thresholds, age gating technology should be correct both near to and far from the thresholds. Nevertheless, proponents of age gating by estimation might contend that offline age gating is often equally fuzzy around the thresholds and that this works well enough for most purposes. Ultimately the determination of "well enough" is a normative and political question that may vary by communities, states, and governance goals.

While some companies that provide age estimation based on face images are willing to participate NIST's FATE-AEV program, few provide open-source code or training data that would allow audits by academic researchers or humans rights advocates. This makes it impossible to tell to what extent self-reported or NIST correctness metrics actually correspond to performance on the wide variety of demographics, image

[125] *Id.*

variations, and other real-world complications that challenge deployed age estimators. This lack of transparency is essential to consider when reasoning about the burdens imposed by age gating using face images, since this type of age gating raises significant concerns about scope creep and data misuse. Face images are highly identifiable, making their collection and use for identity verification, profiling, tracking, and surveillance especially incentivized.[126] Even when commercial age gating services make promises about how they handle face images to protect user privacy, these claims are often a far cry from true data minimization or unlinkability and generally impossible to independently verify.

Voice-based approaches to age estimation[127] use recorded vocal patterns and machine learning models in order to estimate the age of the speaker. The accuracy of these approaches varies depending on whether the model attempts to estimate the exact age of the speaker or only whether the speaker was a member of a particular age group (e.g., "seniors"). The most optimistic results distinguishing senior speakers from non-senior speakers have published accuracies in the mid-90% range;[128] however, closed world hypotheses (training and testing the machine learning algorithm on fixed set of example voices) and long recording lengths (often entire conversations) raise questions about the correctness limitations and potential interaction costs of voice-based inference. Additionally, automatic speaker recognition systems are known to be vulnerable to voice disguises,[129]

---

[126] *See* Sharon Naker & Dov Greenbaum. *Now You See Me: Now You Still Do: Facial Recognition Technology and the Growing Lack of Privacy*. 23 BUJ SCI. & TECH. L. 88 (2017).

[127] Florian Metze et al., *Comparison of Four Approaches to Age and Gender Recognition for Telephone Applications*, 2007 IEEE INT'L CONF. ON ACOUSTICS, SPEECH & SIGNAL PROCESSING (2007); Nobuaki Minematsu et al., *Automatic Estimation of One's Age with His/Her Speech Based Upon Acoustic Modeling Techniques of Speakers*, 2002 IEEE INT'L CONF. ON ACOUSTICS, SPEECH, & SIGNAL PROCESSING (2002); Tobias Bocklet et al., *Age and Gender Recognition for Telephone Applications Based on GMM Supervectors and Support Vector Machines,* 2008 IEEE INT'L CONF. ON ACOUSTICS, SPEECH AND SIGNAL PROCESSING 2644 (2008).

[128] Nobuaki Minematsu et al., *supra* note 127.

[129] Mireia Farrús, *Voice Disguise in Automatic Speaker Recognition*, ACM Comput. Surv. 51(4):68 (2018).

suggesting that voice-based age recognition systems may likewise be easily circumvented. At the very least, much more research would be needed before voice-based age gating could be deployed effectively.

Other biometric age estimation approaches include the use of gait information from photos or videos of users showing movement of the arms and legs.[130] These approaches use a variety of data processing techniques and machine learning classification methods. Similar to approaches using other biometrics, these approaches are still under active research and pose correctness and interaction cost concerns for age gating applications.

Importantly, all forms of biometric age estimation require the use of data that can support many types of inferences in addition to—and often more accurately than—age estimation. Images of users' faces can be used to discriminate based on race, gender, or other legally protected characteristics. Voice recordings could be used with generative artificial intelligence models to create new audio of a user saying something they never actually said. Video footage of users' gaits could indicate various medical conditions that users would prefer not to share. This raises significant privacy objections to the unconsidered deployment of biometric methods for age gating and places a strong priority on unlinkability and data minimization for any age gating technologies involving biometrics. Source code and data transparency, regular audits, and regulatory oversight would be necessary (but perhaps not sufficient) to enforce that biometric data be used only for the purpose of determining age for governance thresholds and not stored, linked, transmitted, or used for any other purpose.

---

[130] Chi Xu et al., *Uncertainty-Aware Gait-Based Age Estimation and Its Applications*, 3 IEEE TRANS. ON BIOMETRICS, BEHAVIOR & IDENTITY SCI. 479 (2021); Chi Xu et al., *Real-Time Gait-Based Age Estimation and Gender Classification from a Single Image*, 2021 IEEE WINTER CONFERENCE ON APPLICATIONS OF COMPUTER VISION (WACV) 3459 (2021); Haiping Zhu et al., *Ordinal Distribution Regression for Gait-Based Age Estimation*, 63 SCI. CHINA INF. SCI. 1 (2020); Jiwen Lu & Yap-Peng Tan, *Gait-Based Human Age Estimation*, 5 IEEE Trans. Info. Forensics & Sec. 761 (2010); Shaoxiong Zhang et al., *Gait-Based Age Estimation with Deep Convolutional Neural Network*, 2019 IEEE INT'L CONF. ON BIOMETRICS (ICB) (2019).

Furthermore, many members of the public are rightly concerned that the collection of biometric data, especially face images, could subject them to unwanted harms.[131] Even with technical guarantees of data minimization or unlinkability, many users would be uncomfortable providing biometric data to engage with online services. This could cause chilling effects with First Amendment implications if users elect not to provide data for age estimation and are inappropriately blocked from online experiences.

2. Capacity and Knowledge Testing

Capacity and knowledge testing techniques attempt to estimate user age based on capacities, such as ability to solve puzzles or language tests, or based on possession of specific knowledge, such as generational trends that might be less salient to users from other generations. These approaches are less well-studied than biometric age estimation.

Hecht et al. used an acoustic Gaussian mixture model and support vector machine system (GMMSVM) combined with a word-based recognition system to estimate individuals' ages based on both their vocal patterns (biometric) and their word use patterns (capacity).[132] Despite the novelty of this approach, the error was still over 20%, far too high for practical age gating applications.

Dukellis & Butcher proposed a form of online advertisement that attempts to estimate user age based on their response to questions with answers that "are typically known only to someone that meets a particular threshold age or is in a specific age range…For example, for a user who claims to be born in 1977 in the United States, a cartoon from 1985 can be shown and a question posed as to who the cartoon character is."[133] However, they did not test the accuracy of this approach,

---

[131] Shikun Zhang, Yuanyuan Feng & Norman Sadeh, *Facial Recognition: Understanding Privacy Concerns and Attitudes Across Increasingly Diverse Deployment Scenarios*, PROCEEDINGS OF THE SEVENTEENTH SYMPOSIUM ON USABLE PRIVACY AND SECURITY 243 (2021).

[132] Ron M. Hecht et al., *Age Verification Using a Hybrid Speech Processing Approach*, INTERSPEECH 184 (2009).

[133] John Nicholas Dukellis & Derek Butcher, *Ads That Verify User Age* 2, 4 TECHNICAL DISCLOSURE COMMONS (2019).

which seems especially prone to users looking up answers online to circumvent the test.

Overall, capacity and knowledge testing seem unlikely to produce age estimates that are both correct and resistant to user circumvention.

3. Profiling

Profiling techniques attempt to estimate user age based on online behavior profiles, such as "how long they spend on a webpage, where their cursor hovers, the times of day they access services and their interests, location, and friends."[134]

This type of age estimation is becoming more common: In 2025, YouTube announced that it would estimate user ages based on data such as "the types of videos a user is searching for, the categories of videos they have watched, or the longevity of the account."[135] Shortly thereafter, OpenAI announced that it would attempt to identify underage users based on their interactions with ChatGPT.[136] In the academic literature, Hossain & Haberfeld used tapping patterns on touchscreens to infer whether a user is below an age threshold with approximately 73% accuracy on phones and 82% accuracy on tablets.[137]

Age estimation based on online behavior profiles could be more difficult for users to circumvent but requires extended data collection about user behavior that threatens unlinkability, increases the chance of scope creep and data misuse, and raises similar privacy concerns as all other forms of online tracking.

C. Age Verification

Unlike self-attestation and age estimation, age verification technologies use a trusted information source or authoritative verification service to confirm a user's age. Figure 2 depicts the high-level schematics of many proposed age verification systems. Two, often decoupled, processes involve the user, the

---

[134] Adib, et al., *supra* note 117.
[135] James Besar, *supra* note 82.
[136] OpenAI, *supra* note 82.
[137] Md Shafaeat Hossain & Carl Haberfeld, *Touch Behavior Based Age Estimation Toward Enhancing Child Safety*, 2020 IEEE INT'L JOINT CONF. ON BIOMETRICS (IJCB) (2020).

first party deploying the age gate, and a trusted third party. The user starts by obtaining a verifiable certification from the third party that they are a certain age or meet a certain age threshold. This process may be mediated by the first party to facilitate usability, but does not necessarily involve the first party. The user then provides this certification to the first party, who is able to verify its authenticity and apply the correct governance rules to the user. Specific variations on this general process have taken many forms, including traditional IDs, digital IDs, anonymous credentialing systems, and versions of passwordless authentication.

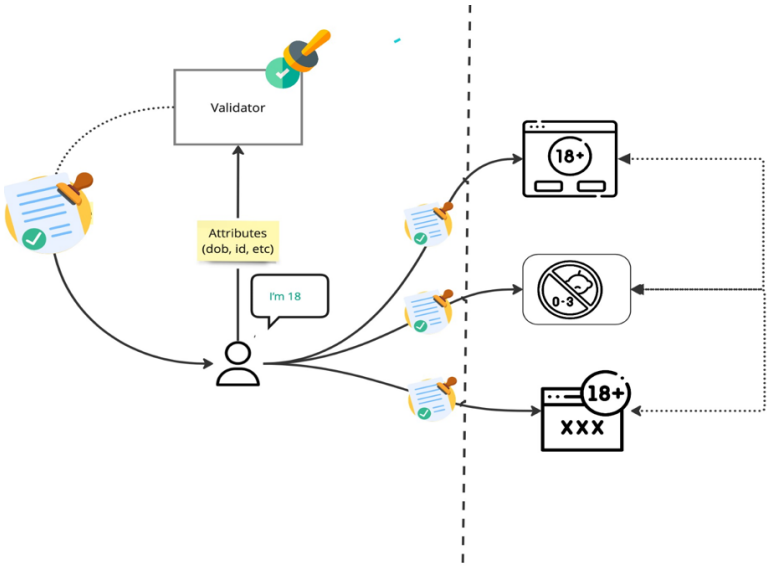

Figure 2: General schematic of many age verification systems. The user obtains a verifiable certification of age from a third-party validator then provides this certification to first-parties, who can verify its authenticity and apply the correct governance rule to the user.

1. Traditional IDs

Many online platforms use government-issued identification (e.g., driving licenses or passports) or credit cards as a verifiable certification that the holder of the ID meets an age threshold.[138] The user enters identifying information from

[138] *E.g.*, Google, *Access Age-Restricted Content & Features*, GOOGLE ACCOUNT HELP, https://support.google.com/accounts/answer/10071085?hl=en (last visited Jul. 1, 2025) [https://perma.cc/LDL8-HDJ7]; BankCard USA, *Age*

the ID or a photograph of the ID itself, and the age gating software automatically consults with the Application Program Interface (API) of the trusted third-party (e.g., the state ID database or the credit card company). If the third-party confirms that the ID is valid and the user meets the age threshold, the first-party service allows the user through the gate.

Since these forms of identification are the most commonly used for offline age gating, they are familiar, readily available, and contain various anti-counterfeiting features of varying usefulness in the online context. However, the use of traditional identification for online age gating suffers from an important privacy risk: the first party can learn more information about the individual than is required for age gating, as the identification documents contain many personal details. This makes age gating based on traditional IDs especially prone to scope creep and data misuse. Information on traditional IDs can be readily linked to user profiles unrelated to age-based governance or used to perform identity verification and track users across the Internet. Designing and/or regulating age verification systems based on traditional IDs that provide trustworthy unlinkability and data minimization is a challenging design prospect that may be infeasible. Age gating systems that simply require users to submit images of or numbers from traditional IDs should therefore be avoided.

2. Digital IDs

Many state and federal governments are currently interested in the development of digital identity documents (digital IDs), also known as electronic identity documents (eIDs), that could replace or supplement traditional IDs. These eIDs can easily include age information and could potentially be used for online age gating.

The International Civil Aviation Organization (ICAO) has standardized identity documents that combine traditional documentation with electronic chips that can store and execute

---

Verification, https://www.bankcardusa.com/online-age-verification/ (last visited Jul. 1, 2025) [https://perma.cc/P5SJ-54R4].

cryptographic primitives.[139] ePassports implementing this standard are being widely adopted and issued in at least 140 states around the world with the backing of the United Nations and European Union.[140] These eIDs can store several attributes, such as the birth date and full name of the carrying individual. An age-gating service could query a user's ePassport as to whether the cardholder's age is above a certain date.[141] The information on the ePassport is cryptographically signed by the issuing country, allowing the service to verify the authenticity of the data.

Digital IDs are considered part of digital public infrastructure (DPI) around the world, including India's Aadhaar system, China's Cyber Trusted Identity (CTID) and Real-Name Decentralized Identifier (RealDID) systems, Singapore's National Digital Identity (NDI) system, and several systems in Nordic countries. The International Organization for Standardization has defined specifications for mobile driving licenses,[142] and the World Bank's Identification for Development (ID4D) advocates for the widespread deployment of eIDs as a part of its sustainable development goals.[143] However, digital ID systems, including those promoted by ID4D, have raised human rights concerns due to loss of privacy and the potential for "exclusion from public and private services...biometric exclusion, discrimination along new cleavages, and the many harms associated with surveillance capitalism."[144]

---

[139] Andreas Poller et al., *Electronic Identity Cards for User Authentication - Promise and Practice*, 10 IEEE SECURITY & PRIVACY MAG. 46 (2012).

[140] ICAO, *ePassport Basics*, https://www.icao.int/Security/FAL/PKD/Pages/ePassport-Basics.aspx (last visited May 2, 2024) [https://perma.cc/AF39-8RXK].

[141] Ingo Naumann & Giles Hogben, *Privacy Features of European eID Card Specifications*, 2008.8 NETWORK SECURITY 9 (2008).

[142] International Organization for Standardization, *ISO/IEC 18013-5:2021 Personal Identification — ISO-Compliant Driving License* (2021).

[143] WORLD BANK, *Identification for Development (ID4D)*, https://id4d.worldbank.org (last visited Apr. 15, 2025) [https://perma.cc/4XLJ-LA8Z].

[144] CENTER FOR HUMAN RIGHTS AND GLOBAL JUSTICE, N.Y.U., *Paving a Digital Road to Hell? A Primer on the Role of the World Bank and Global Networks in Promoting Digital ID*, (2022), https://chrgj.org/2022-06-paving-

In the United States, several states have issued or are considering mobile driver's licenses (mDL) that act as digital IDs. Many of these mDLs work via an app provided by the state, and some support mobile wallets (e.g., Apple Wallet, Google Wallet, or Samsung Wallet). The U.S. Transportation Security Administration (TSA) has begun accepting mDLs from selected states as valid IDs for airport checkpoints. The increasing availability of digital IDs in the United States has also raised concerns about privacy, personal liberties, and human rights. The American Civil Liberties Union (ACLU) has published state legislative recommendations regarding digital IDs addressing these issues and other concerns.[145]

For the purposes of age gating, digital IDs pose most of the same problems as traditional IDs, while introducing even more concerns about scope creep and data misuse. The creation of an Internet infrastructure around digital IDs makes it more likely that age gating becomes supplanted by identity verification, in which online services decide that it is easier to collect users' complete digital ID information instead of just their age. This is further incentivized by profit incentives for user profiling and tracking and government incentives toward surveillance. In the worst case, digital IDs could become a "super cookie" by which user activity across many online services is able to be linked.[146] Age gating that requires users to submit information from a digital ID could contribute to this outcome.

These drawbacks clearly indicate that digital IDs should not be used naïvely for online age gating. France's National Commission on Informatics and Liberty (CNIL) examined the privacy risks of existing age verification techniques and recommended a design that relies on a trusted independent third-party that would be regularly audited and certified by a

---

digital-road-to-hell/ (last visited Apr. 16, 2025) [https://perma.cc/5A5H-9T5A].

[145] AMERICAN CIVIL LIBERTIES UNION, *ACLU Digital ID State Legislative Recommendations* (Oct. 2024), https://assets.aclu.org/live/uploads/2024/10/ACLU-Digital-ID-State-Legislative-Recommendations-version-1.0-October-2024-1.pdf [https://perma.cc/8X8W-CEMS].

[146] *Id.*

reputable agency.[147] The ACLU similarly advocates for a digital ID system that is private, transparent, and provides unlinkable verification.[148] The following sections describe progress toward this ideal.

3. Anonymous Credentials

In the 1980s, Chaum introduced the notion of designing anonymous credential (AC) authentication systems that allow individuals to prove one or more "attributes," such as age, to potentially untrustworthy institutions.[149] Several AC schemes have since been implemented using cryptographic primitives, such as public key cryptography,[150] zero-knowledge proofs,[151] and blind digital signatures.[152]

Baldimtsi & Lysyanskaya proposed using blind signatures to certify attributes such as age.[153] In this scheme, age verification is performed by a protocol between a user, one or

---

[147] COMMISSION NATIONALE DE L'INFORMATIQUE ET DES LIBERTÉS (CNIL), *Online Age Verification: Balancing Privacy and the Protection of Minors* (Sept. 22, 2022), https://www.cnil.fr/en/online-age-verification-balancing-privacy-and-protection-minors (last visited Apr. 5, 2024) [https://perma.cc/6QQV-A8E7].

[148] AMERICAN CIVIL LIBERTIES UNION, *supra* note 145, 3.

[149] David Chaum, *Security Without Identification: Transaction Systems to Make Big Brother Obsolete*, 28 COMMUNICATIONS OF THE ACM 1030 (Oct. 1985).

[150] Public key cryptography uses a pair of mathematically linked keys: a public key, which is shared openly, and a private key, which is kept secret. These keys can be used for mathematically linked operations, such as encrypting and decrypting messages or creating and verifying signatures.

[151] A zero-knowledge proof is a cryptographic process by which one party can prove to another party that a given statement is true, without revealing any additional information.

[152] Jan Camenisch & Anna Lysyanskaya, *An Efficient System for Non-Transferable Anonymous Credentials with Optional Anonymity Revocation*, ADVANCES IN CRYPTOLOGY—EUROCRYPT 2001: INT'L CONF. ON THEORY & APPLICATION CRYPTOGRAPHIC TECHS. 93 (2001); CHRISTIAN PAQUIN & GREG ZAVERUCHA, MICROSOFT CORP., *U-Prove Cryptographic Specification V1.1, Revision 3* (Dec. 2013), https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Cryptographic20Specification20V1.1.pdf [https://perma.cc/AC5P-PNCU].

[153] Foteini Baldimtsi & Anna Lysyanskaya, *Anonymous Credentials Light*, PROCEEDINGS OF THE 2013 ACM SIGSAC CONF. ON COMPUT. & COMMC'NS SECURITY 1087 (2013).

more untrusted age-gated services, and an identity verifier (IV) trusted by all parties as follows: The user generates a cryptographic commitment, C, which binds together their age (or another relevant attribute) and a secret key they possess. The user sends C to the IV, which performs the blind signature with attributes protocol: First, the IV verifies that the user actually possesses the secret key by performing a cryptographic check against the user's public key. Second, the IV verifies the user's age claimed in C through some external sociotechnical process. If both verifications succeed, the IV creates a new unlinkable commitment, C', binding together the user's secret key and their age. The IV signs C' using its own secret key, producing digital signature σ. The IV finally sends C' and σ back to the user. The user can then send C' and σ to an untrusted age-gated service in order to prove their age without revealing any additional information. The service can verify whether the age bound by C' meets the threshold set by the age gate and confirm that C' was issued by the identity verifier (IV) by verifying σ. This process can be readily modified to include additional attributes to make the credential produced by IV single-use, so the user can only use C',σ to verify their age once before they must return to IV for another credential, or to place an expiration date on the use of the credential.

This AC process provides several beneficial properties. The age-gated service can be confident that the age is correct, because it trusts the age verification process performed by the IV. If the user attempts to modify a commitment C' after it has been signed by the IV, e.g., to increase the claimed age, the digital signature σ will no longer verify. If the user attempts to create C',σ themselves from scratch, the service will be able to detect that these do not come from the IV. If the user transfers C',σ to another user, the secret key claimed in C' will not verify, so the transfer can be detected and rejected by the age-gated service. The IV does not learn what age-gated service the user is trying to access, and the age-gated service learns nothing about the user other than their age, providing effective data minimization. Finally, the protocol is unlinkable under standard cryptographic assumptions.[154]

[154] *Id.*

The above AC process is just one specific example of the many AC processes that the cryptography community has been developing for many years in order to achieve different security properties, to be practical in various use cases, [155] or to align with various regulations[156] or standards.[157] While these processes show great promise, there are still practical concerns that must be considered.

AC systems push the job of age determination onto a trusted identity verifier. This leaves open the questions of what entity acts as this verifier, why are they trusted, whether they are incentivized to provide this IV service, and how they determine a user's age. Obvious candidates for IV providers include organizations that already have enough information about users to verify their identity and age, especially state and federal governments. Phone plan providers, operating system vendors, banks, and credit card companies could also potentially provide IV services. Each of these entities may be trusted by some users and some age-gated services but not others. An ecosystem could develop in which multiple entities provide IV services, allowing users to choose which entity they prefer and providing competition as incentive for these entities to limit prices and avoid misbehavior (such as inappropriately including more information about the user in commitments sent to age-gated services or by requiring details of what age-gated content or actions users seek to access in order to provide IV support). The risk of misbehaving IVs could also be mediated by regular audits and the potential for reputational damage, similar to that faced by misbehaving certificate authorities in the current public key infrastructure.[158]

---

[155] Saqib A. Kakvi et al., *SOK: Anonymous Credentials,* SECURITY STANDARDISATION RESEARCH 8TH INTERNATIONAL CONFERENCE (SSR 2023) *reprinted in* 13895 LECTURE NOTES IN COMPUTER SCIENCE (LNCS) 129 (Felix Günther & Julia Hesse, eds., 2023);

[156] Nicolas Desmoulins et al., *Making BBS Anonymous Credentials eIDAS 2.0 Compliant,* CRYPTOLOGY EPRINT ARCHIVE (2025).

[157] Matteo Frigo & Abhi Shelat. *Anonymous credentials from ECDSA,* CRYPTOLOGY EPRINT ARCHIVE (2024).

[158] Michael Alan Specter, *The Economics of Cryptographic Trust: Understanding Certificate Authorities* (2016) (S.M. thesis, Massachusetts Institute of Technology).

AC protocols also leave room for certain types of user circumvention. These protocols generally assume that each user will have a unique pair of public and secret keys, and that they will not share their secret key with others under any circumstance. In practice, secret keys are typically stored on user devices (laptops, smartphones, etc.) by key management software and may be bound to a hardware token. Users often let others borrow their devices, e.g., a child borrowing a parent's laptop or phone. Without alternative methods for binding keys to users, the risk of circumvention by device sharing would have to be addressed by social norms, e.g., when it is appropriate to allow a child to use an adult's device unsupervised.

Most users are also not technically savvy enough to manage multi-party cryptographic protocols manually. This means that the AC process would have to be mediated by software running on the user's device. This software would have to be trusted, otherwise it could break unlinkability and send information about the age-gated service to the IV or identifying information about the user to the age-gated service.

Importantly, none of these concerns imply that AC systems are fundamentally infeasible or unrecommended for online age gating. A world in which all online age gating, including the many self-attestation gates and traditional ID gates currently present on the Internet, were converted to trusted AC-based gates would provide a great deal more privacy for users and effectiveness in enforcing age-based governance. With sufficient software development effort, the user experience could be nearly seamless. Software on the user's device could handle the acquisition and storage of credentials from the IV and the transfer of credentials to age gates. Recent efforts to make anonymous credentialing more practical and effective have incorporated innovations in passwordless authentication.

4. Passwordless Authentication

The Fast IDentity Online (FIDO) Alliance association and the World Wide Web Consortium (W3C) lead the FIDO2 Project to facilitate new forms of passwordless authentication

for the web.[159] The collection of authentication standards they have produced incorporates public key cryptography with biometrics, voice, and facial recognition.

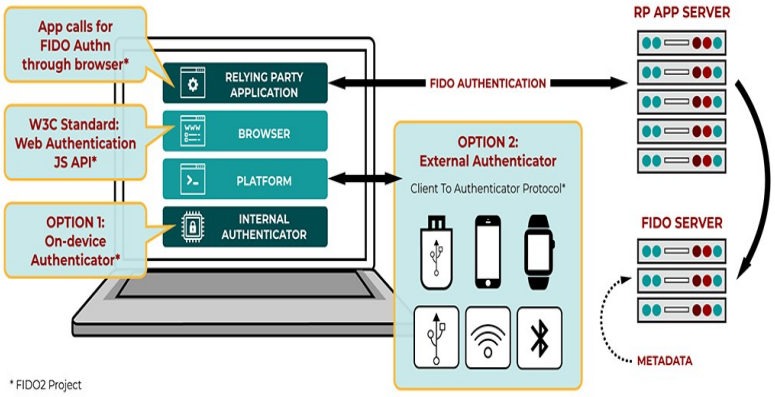

Figure 3: FIDO2 system diagram. Figure by FIDO Alliance.

The FIDO2 authentication mechanism works in three main phases (ceremonies) involving a user, an authenticator (such as a physical security key, trusted platform module (TPM), embedded secure element (eSE), or a device with face, fingerprint, or other biometric ID), and the relying party (service) to which the user seeks to authenticate (Figure 3[160]). First, the user registers with the relying party to generate a private/public key. The private key is kept by the authenticator and the public key is kept by the relying party. When a user attempts to authenticate to the service later, the relying party sends a cryptographic challenge to the user. The authenticator replies to the challenge using the saved key, which is verified by the relying party using the public key.

In addition to security properties, the FIDO protocols attempt to balance privacy concerns. The FIDO Privacy Principles[161] include informed consent; multiple, cross-context

---

[159] FIDO ALLIANCE, *FIDO Authentication: A Passwordless Vision*, https://fidoalliance.org/fido2/ (last visited Apr. 15, 2024) [https://perma.cc/5G63-UVVZ].

[160] Illustration of a FIDO workflow, in FIDO ALLIANCE, *User Authentication Specifications Overview*, fig. https://fidoalliance.org/specifications/ (last visited Aug. 1, 2024) [https://perma.cc/XJC3-H5TA].

[161] FIDO ALLIANCE, *FIDO Privacy: FIDO Alliance White Paper* (Jan. 19, 2016), https://fidoalliance.org/wp-

identities; limited collection of personal information; limited use of users' data for FIDO-related purposes such as registration, user verification, and authentication; prevention of user identification outside the FIDO context; keeping biometric data on users' computer; and protection of unauthorized access to FIDO-related information.[162]

Yeoh et al. have proposed combining the FIDO2 protocol with autonomous credentialing and eIDs to facilitate unlinkable sharing of attributes, such as age, with the relying party using a non-interactive zero-knowledge proof system.[163] This FIDO-AC protocol combines the unlinkability and data minimization properties of anonymous credentialing with the security of FIDO2 and the convenience of digital IDs. FIDO-AC and related ideas are still under active research, but they show promise for age gating and other forms of attribute-based online governance.

V. Synthesis & Recommendations

The current landscape of online age gating is rapidly evolving and hotly contested. Unfortunately, participants in debates about online age gating often conflate descriptive claims about what may be technologically feasible with normative claims about what may be desirable, socially appropriate, and legitimate. Making matters worse, participants often fail to appreciate how both sets of claims vary considerably by sociotechnical context and change dynamically over time, along with, and often driven by, changes in technology and social conditions.

To effectively reconcile these complex, interdisciplinary issues, we must consider both the social context in which age gating may be deployed as well as the affordances, burdens, and tradeoffs associated with specific technical

---

content/uploads/FIDO__Privacy_White_Paper_Jan_2016.pdf [https://perma.cc/Q3GG-G33K].

[162] FIDO ALLIANCE, *User Authentication Specifications Overview*, https://fidoalliance.org/specifications/ (last visited Aug. 1, 2024) [https://perma.cc/EW9Z-MAAR].

[163] Wei-Zhu Yeoh et al., *Fast IDentity Online with Anonymous Credentials (FIDO-AC)*, PROCEEDINGS OF THE 32ND USENIX SECURITY SYMPOSIUM, 3029 (2023).

implementations in appropriately characterized sociotechnical environments.

We make the following recommendations for legislatures, regulatory agencies, and courts looking for guidance on reasoning about online age gating.

First, recognize that for many reasons, offline age gating is widespread, reasonably effective, and minimally controversial. This observation provides a baseline, a reference set of burdens and affordances already deemed broadly acceptable, both legally and normatively. It is important to recognize various social and technical contingencies baked into this norm. In some contexts, the burdens of offline age gating may fall disproportionately on some groups (e.g., those who lack access to required documentation) and mitigating those concerns remains a challenging and important task. This possibility does not, in general, undermine the case for offline age gating; rather, it highlights the context-specific nature of evaluating and designing age gating within broader systems of governance and the background sociotechnical environment (e.g., concerning access to legal documentation and available alternatives).

Second, recognize that the transposition of age gating from offline to online environments changes the sociotechnical context, including the actors, technologies, information flows, and relationships, and consequently also the burdens and affordances of the corresponding age-associated governance. New social costs and burdens, such as the possibility of linking information from the age gating process to other details about a user, the potential chilling of speech in otherwise anonymous or pseudonymous fora, the fact that data selling is the norm online, and the potential decrease in usability and accessibility, may overwhelm the intended benefits of the governance. Conversely, new affordances, such as the possibility of provably limiting circumvention, linking, and/or profiling, as well as new potential to protect vulnerable individuals, may make some new costs and burdens acceptable. These observations lead us to emphasize the sociotechnical and contextual contingency of evaluating and designing age gating systems within broader systems of governance.

Although significant progress has been made in the development of online age gating technologies, none (yet) provide the same affordances as offline age gating without additional complications and burdens for stakeholders. For example, online self-attestation, while widespread, is vulnerable to circumvention and little more than theater. Current age estimation approaches based on biometrics or behavior profiles have limited accuracy, limited transparency, and/or raise serious privacy concerns. Age verification techniques based on trusted information sources and cryptography techniques are the most promising, but additional work is still needed to reduce hardware requirements, stress test privacy claims, and conduct user studies to determine usability impacts.

The relevance and weight of the different burdens we have identified vary across age gating methods as well as the sociotechnical contexts in which such methods may be deployed. For example, in contexts raising First Amendment challenges, whether and how to weigh privacy burdens depends, we suggest, on the actors involved, context-specific privacy norms, and the manner in which perceived privacy violations *actually* impact speech related activities. Too often, courts and scholars presume (or accept speculative assertions about) chilling effects or other related burdens on speech without empirical support or appreciation of how the sociotechnical environment shapes the relationship between privacy and speech.

Nevertheless, age gating technologies are rapidly changing, and further developments may soon yield an effective age determination method with minimal tradeoffs. Should this occur, we hope that case law and social norms will evolve to reflect the new technological paradigm. Ascertaining changes in social norms and their implications can be challenging, and could be guided by rigorous interdisciplinary frameworks, such

as contextual integrity,[164] governing knowledge commons,[165] or a combination of both.[166]

Third, recognize that there are several valid *content-neutral* motivations for online age gating. While content-based age gating still comprises the majority of case law and policymaking on the topic, we specifically highlight capability-based motivations and vulnerability-based motivations as reasons for age gating independent of the content involved. It is important to avoid overreaching precedent blocking (or mandating) online age gating for content or free speech motivations that might limit the potential of capability-based or vulnerability-based governance.

A. Free Speech Coalition, Inc. v. Paxton

Note to readers: We wrote this article prior to oral arguments in the Supreme Court and posted a draft on ssrn.com. The Electronic Privacy Information Center subsequently cited our article in its amicus brief.[167] This section originally continued our discussion of the Fifth Circuit opinions and made modest recommendations to the Supreme Court. We now provide our original discussion followed by a brief update and reflection based on the Court's decision.

Recall from our earlier discussion that the Fifth Circuit majority reasoned that at least one of three methods of online age verification, "(1) government ID, (2) facial appearance, or (3) some other available information used to infer the user's

---

[164] HELEN NISSENBAUM, PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE (2010); Helen Nissenbaum, *Respect for Context as a Benchmark for Privacy Online: What It Is and Isn't*, CAHIER DE PROSPECTIVE: THE FUTURES OF PRIVACY 19 (Carine Dartiguepeyrou, ed., 2014).

[165] GOVERNING KNOWLEDGE COMMONS (Brett M. Frischmann, Michael J. Madison, & Katherine Strandburg, eds., 2014).

[166] Yan Shvartzshnaider et al., *GKC-CI: A Unifying Framework for Contextual Norms and Information Governance*, 73.9 J. ASS'N FOR INFO. SCI. & TECH. 1297 (2022).

[167] *See* Br. of Amicus Curiae Electronic Privacy Information Center in Support of Neither Party, *Free Speech Coal., Inc. v. Paxton*, No. 23-1122 (September 23, 2024), at https://epic.org/wp-content/uploads/2024/09/EPIC-Amicus-FSC-v.-Paxton.pdf [https://perma.cc/ZS6X-6DYN].

age," would "have no more impact on privacy than … in-person age verification."[168] As we now explain, this factual claim is simply incorrect.

The Fifth Circuit majority appropriately reasoned about the varying burdens imposed by offline versus online age gating, but the claim of no additional privacy impact is insufficiently supported. The age determination methods referenced by the majority correspond to methods we evaluated in this article, namely, age estimation with biometrics (face appearance) and age verification using traditional IDs. We discussed how current implementations of these methods, while ongoing research areas, are still not equivalent to offline methods in terms of the properties they provide and the threat models to which they are vulnerable. Showing your face or a government identification offline—in person, in a particular geographic location and social setting to particular people—is simply not the same as undergoing automated facial recognition or submitting a government identification online. User devices, cloud servers, and network infrastructure run opaque software programs developed by first- and third-party organizations often with strong incentives to record (perfectly, and often in perpetuity), link, profile, and monetize user information. Although the same types of information are used in both offline and online scenarios, the technological means, sociotechnical context, and resulting *burdens* all differ dramatically and, in a manner, are relevant to legal analysis.

The Fifth Circuit majority further explained that "even were there a gap in privacy concerns as large as plaintiffs suggest, we decline to adopt their notion that such a gap matters. In short, no binding precedent compels us to depart from *Ginsberg* on privacy grounds, and we decline to do so."[169] This legal argument about the irrelevance of privacy burdens seems to bring the decision outside the purview of the types of empirical (factual) arguments based on technological foundations that we have advanced. We are not persuaded, however, for three interrelated reasons.

---

[168] *Free Speech Coal.*, 95 F.4th at 271.
[169] *Id.*

First, with respect to binding precedent, when the Supreme Court struck down sections of the CDA in *Reno*, the Court considered the burdens of age verification relevant, albeit from a different angle. The Court evaluated whether the age verification defense sufficiently reduced the burdens on adult speech; it did not frame its analysis in terms of additional burdens posed by having to submit to privacy-invasive age verification. Yet the notion that such privacy invasions might chill or otherwise burden adult speech was implicit in the Court's reasoning. Had the available means for online age verification functioned identically to those present and widely used offline, as in *Ginsberg*, for example, then the age verification defense in the CDA arguably would have been (seen by the Court) as much more effective in reducing the burdens on adult speech.[170]

Second, the Fifth Circuit *Paxton* majority fails to appreciate the relevance of the baseline technological facts that shape privacy and speech burdens in the context of *Ginsberg* itself (as discussed in Section I and Section II).

Third, privacy and speech (burdens) are often interdependent. Unless one woodenly insists upon their independence (contrary to reality as well as Supreme Court jurisprudence[171]), it is hard to see how a substantial privacy gap is not likely to be associated with speech burdens that matter. Simply put, if privacy burdens deter speech activities, then the burdens should be considered relevant. Anonymity (even perceived anonymity that is really pseudonymity) impacts what people are willing to say, listen to, watch, and do online. It is often just assumed that a lack of privacy chills speech, but we think it raises important empirical questions that matter, which is why we argue that a privacy gap/burden is relevant but not dispositive. A more nuanced approach to the case would have

---

[170] Oddly, the Supreme Court in *Paxton* ignores any such burdens (or perhaps slightly more charitably, presumes them to be merely incidental) and instead characterizes the age verification defense in the CDA as "illusory" because "existing technology" was ineffective. 606 U.S. 461 at 23. Yet, as we discuss below, the Court in *Paxton* never evaluates, in a meaningful, empirical sense, modern age verification technologies in terms of burdens or effectiveness.

[171] *See, e.g.*, *McIntyre v. Ohio Elections Commission*, 514 U.S. 334 (1995).

meaningfully engaged the actual tradeoffs involved in mandating age gating given current technological methods. The majority's attempt to frame the legal argument in terms of binding precedent and a forced or chosen departure from *Ginsberg* also seems strained and heavily reliant on the factual premises about age verification technologies.

The U.S. Supreme Court granted certiorari in this case.[172] In our view, the Supreme Court should require serious engagement with the technological details and corresponding tradeoffs. It should reject (and thus refuse to countenance) speculative, hyperbolic, and anecdotal assertions about technical details, implementations, burdens, and chilling effects. Further, it should take a moderate stance and avoid overreaching precedent. On one hand, precedent blocking most or all forms of online age gating on free speech grounds could limit content-neutral age gating laws and the deployment of less burdensome age gating technologies if and when they are developed, leaving minors and other potentially vulnerable individuals without appropriate governance protections. On the other hand, precedent supporting or enforcing content-based online age gating using current technologies without sufficient consideration and oversight could expose users of all ages to substantial burdens, including personal data collection that would further expand the data-driven "surveillance capitalism" of the modern Internet.[173]

What did the Supreme Court do? Did the Court follow our advice? Did it chart a different path?

The Supreme Court held that the Texas age gating law triggers, and survives, review under intermediate scrutiny because it only incidentally burdens the protected speech of adults. The Court thus rejected arguments from both sides, and followed by lower courts, to apply either strict scrutiny or rational basis. We are not bothered by the decision to apply intermediate scrutiny, as strict scrutiny and rational basis are too often outcome-determinative and fail to support or engage in the more nuanced, interdisciplinary analysis we suggest is

---

[172] *Free Speech Coal., Inc. v. Paxton*, 95 F 4th 263, *cert. granted*, 144 S.Ct. 2714 (Mem) (U.S. July 2, 2024) (No. 23-1112).

[173] *See* SHOSHANA ZUBOFF, *supra* note 87.

necessary.[174] Yet, as we explain below, we are seriously troubled by the Court's reasoning and application of intermediate scrutiny, especially its deference to legislative decision-making and failure to honestly engage the technological facts of online age gating and corresponding privacy and speech burdens. Notably, in her dissent, joined by Justices Sotomayor and Jackson, Justice Kagan engages more seriously and honestly with the technological details and corresponding burdens. For example, Justice Kagan emphasizes how age verification online is nothing like showing an ID offline. She explains how the privacy concerns are quite different and how they can lead to chilling effects and other harms. In this short section, we focus on the majority opinion.[175]

The reasoning employed by the majority to explain and justify the Court's holding rested heavily on the two-tier approach recognized in *Ginsberg* and an unproven and unjustified insistence that age verification offline and online is equivalent, effective, and merely an incidental burden on adult speech. The Court summarizes: "A State may not prohibit adults from accessing content that is obscene only to minors. … But, it may enact laws to prevent minors from accessing such content."[176] The Court then declares the "power to verify age is a necessary component of the power to prevent children's access to [such] content,"[177] which then leads to the following pronouncement: "because the First Amendment permits States to prohibit minors from accessing speech that is obscene to them, it likewise permits States to employ the ordinary and appropriate means of enforcing such a prohibition. Requiring proof of age to access that speech is one such means."[178] The

---

[174] In our view, intermediate scrutiny is defensible so long as age gating technology delivers effective, minimally burdensome, normalized age gating on par with *Ginsberg*, which our analysis in Sections 4 and 5 shows is not (yet) the case.

[175] Our discussion of the Supreme Court decision refers to age verification because that is the terminology used by the Court. As our previous discussion emphasizes, age verification is a sub-category of age gating technologies.

[176] *Id.* at 10.

[177] *Id.* at 14.

[178] *Id.*

Court then discusses how age verification is common offline for various activities (tobacco, firearms, etc.) and obscenity.[179]

The Court never engages in a rigorous examination of age verification online.[180] It never considers the significant differences between offline and online age gating we examined above. The Court acknowledges that technologies have changed dramatically over the past few decades, but it does so only to serve the narrow – results-oriented – conclusion that the need for age gating laws has grown because of Internet expansion, smartphones, and the ease of access to pornography and other content deemed bad for minors.[181] Essentially, the Court superficially considers how changed technological conditions affect the perceived benefits of age verification but deliberately ignores how changed technological conditions affect the burdens. Based on such cursory and myopic analysis, the Court concludes age verification offline and online is "ordinary and appropriate".[182]

More specifically, the Court approves the two methods of age determination appearing in the Texas law: using government-issued identification or transactional (credit card) data.[183] We describe several drawbacks of the naïve use of

---

[179] *Id.* at 14-15.

[180] There is a conclusory paragraph that fails to differentiate among methods or address any burdens. Id. at 33. And there is a similarly conclusory paragraph about privacy concerns and chilling effects. *Id.* at 35.

[181] *Id.* at 26-27.

[182] We admit that it is difficult to say whether the Court deliberately ignores or is just ignorant of the technological facts. The Court makes nonsensical statements like: "Because proof of age performs the same critical function online that it does in person, requiring age verification remains an ordinary and appropriate means of shielding minors in the digital age from material that is obscene to them." *Id.* at 17. The Court later appeals to "common sense" and "centuries of legal tradition" to support the argument that "an age-verification requirement is an ordinary and appropriate means of enforcing an age limit, as is evident both from all other contexts where the law draws lines based on age and from the long, widespread, and unchallenged practice of requiring age verification for in-person sales of material that is obscene to minors." This statement supports the argument we originally made about there not being a per se obstacle to age gating laws. But it is either deliberately ignoring or blissfully ignorant of the fact that age verification online is not the same as age verification for in-person sales. Justice Thomas' effort to ground the majority opinion in "history and tradition" is thus a sham.

[183] *Id.* at 33.

traditional ID-based age verification and digital ID-based age verification in Section IV. Unfortunately, the Court endorses these methods without any consideration of the potential for scope creep and data misuse or any attempt to require verifiable unlinkability or data minimization. Without such safeguards, ID-based age verification risks becoming yet another tool for surveillance, profiling, and tracking, as warned by privacy scholars.[184]

We expect many are worried that online age gating laws will spread and experience function creep in terms of who and what are governed (a political economy concern). Many are also worried that free speech and even the Internet itself are poised for harm or destruction as a result of overzealous online age gating.[185] In the end, we believe such concerns are hyperbolic, at least to the extent they rest on age gating as the underlying cause rather than the many other authoritarian threats to speech, privacy, and an open Internet. The Supreme Court emphasized repeatedly that its analysis focused on, and its holding should be limited to, laws proscribing obscenity to minors.[186] Despite some reasonable concern that state legislatures might define obscenity to minors quite broadly, perhaps extending beyond sexually explicit content, the Court adapted the *Miller* standard to define obscenity for minors as follows: "To be more precise, a State may prevent minors from accessing works that (a) taken as a whole, and under contemporary community standards, appeal to the prurient interest *of minors*; (b) depict or describe specifically defined sexual conduct in a way that is patently offensive *for minors*; and (c) taken as a whole, lack serious literary, artistic, political, or scientific value *for minors*."[187] In theory at least, this legal standard should limit scope creep to other types of speech.

Conclusion

Online age gating is the topic of a heated ongoing debate. Supporters often claim it can protect the vulnerable and govern

---

[184] *See* Sarah Scheffler, *supra* note 92.
[185] *Id.*
[186] *Id.* at 7-10.
[187] *Id.* at 10.

access to certain content. In contrast, those in opposition view the practice as inherent violations of privacy and freedom of speech. In this paper, we examine the legal and policy landscape and review the state of online age gating technologies. We argue that the age gating debate should be refocused around content-neutral capability-based and vulnerability-based motivations. Future consideration of age gating online must take into account the actual technical properties of age gating system designs, including promotion of desired properties and resistance to important threat models. Ultimately, we recognize the complexity of the debate around age gating given the diversity of users, technologies, and ideologies that comprise the Internet. This paper provides a guide to the current state of affairs as well as inspiration for future collaborations between lawyers, policymakers, and technologists on the topic.