The Foreseeable Child:

Age Verification Bypass, the Fraud That Never Was, and the Architecture of Deflected Responsibility

Travis Gilly*

Working draft, July 2026 (v0.1)

Abstract

Age verification is now the law of three continents. Texas and roughly half the American states condition access to adult content or social media on proof of age; the United Kingdom's Online Safety Act requires "highly effective age assurance"; the European Commission's guidelines under Article 28 of the Digital Services Act prescribe age assurance that is accurate, reliable, and robust; Australia bans social media accounts for children under sixteen outright. And children defeat these systems continuously, not with sophisticated tools, but with the ordinary resourcefulness the systems were built to anticipate: a borrowed face, a purchased account, a birthdate field that believes whatever it is told. When a child gets through a mandated gate and is harmed on the other side, a question follows that the entire regulatory architecture is designed never to answer: who is responsible?

This Article maps the architecture of deflection that currently answers "no one," and dismantles it joint by joint. The platform points to compliance with the mandate. The age assurance vendor points to the user's "fraud." The legislature is immune. The blame that remains rolls downhill onto the child who got through. The Article shows that each deflection fails on the deflecting party's own legal materials. The child cannot be the culpable party: no contract containing an age attestation was ever formed between the platform and the child, the operator is legally forbidden from relying on the child's word about her age, and a century of protective-statute doctrine holds that a member of the class a statute protects cannot be charged with defeating it. The bypass cannot be a superseding cause: under the very-hazard rule, the child's foreseeable evasion is the risk the duty exists to address, and it cannot simultaneously be the excuse for breaching that duty. And the compliance shield fails in every jurisdiction once each regime's own text is

*Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation. ORCID: 0009-0007-2954-6313. Correspondence: t.gilly@ai-literacy-labs.org. The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

read: American tort law treats statutory compliance as a floor rather than a ceiling, while the British and European standards define effectiveness partly as resistance to the circumvention methods available to children, which converts every documented child-level bypass from a defense into evidence of noncompliance. The responsibility the architecture deflects lands where the knowledge always sat: on the operators who deployed gates whose failure modes they knew, and on the vendors who certified those gates against error rates they published themselves.

Keywords: age verification; age assurance; minors; superseding cause; negligence per se; contract formation; Online Safety Act; Digital Services Act; Section 230; *Free Speech Coalition v. Paxton*; Restatement (Second) of Torts § 449.

Contents

I. Introduction

Two figures stand at the same gate. The first is eleven years old. The platform's facial estimation system asks her to look into the camera, and she does, wearing her older sister's makeup, and the model, whose published error bands already predicted this outcome for faces near the threshold, waves her through. The second figure is an adult. He did not look into any camera. He paid five dollars on a resale site for an account that someone else's face had already verified into a child's age bracket, and he walked through the gate in the opposite direction, into the child-only spaces the gate was built to seal.[1] Both bypasses were reported within days of the system going live. Neither required a virtual private network, technical skill, or anything a regulator would recognize as sophistication. Both were exactly the events the gate existed to prevent.

Now let harm occur on the far side. The eleven year old is groomed, or sextorted, or bullied into crisis in a space the law said she could not enter. The question every parent, legislator, and plaintiff's lawyer asks next has a strange property: the more parties there are who could answer for the harm, the more confidently each one points at the next. The platform answers that it deployed the verification the statute required. The age assurance vendor answers that its system was defeated by fraud. The legislature that mandated the gate cannot be sued for designing it badly. And the residue of blame, after every institutional actor has exited, settles on the eleven year old, whose makeup is recast as deception, whose gate-crossing is recast as the wrongful act that severed everyone else's responsibility. An industry publication states the resulting doctrine with admirable candor: where a certified, "highly effective" system is bypassed by "a highly sophisticated bad actor," the platform "is generally protected from liability."[2] The

1. Within days of Roblox's January 2026 mandate of facial age verification for chat access, age-verified accounts were reported for sale on eBay for approximately five dollars, permitting a purchaser to enter child-designated chat spaces with a pre-validated minor status; eBay began removing the listings as policy violations. *See Roblox Mandatory Age Verification 2026: Is Your Child Truly Safe?*, Anapol Weiss (May 11, 2026), https://www.anapolweiss.com/blog/roblox-just-changed-everything-what-parents-need-to-know-about-the-new-age-checks/ (last visited July 10, 2026).

2. *Age Assurance 2026: Highly Effective Compliance & ID Solutions*, Klippa (Apr. 22, 2026), https://www.klippa.com/en/blog/information/age-assurance/ (last visited July 10, 2026). The characterization is the industry's; Part VI below tests it against the regulatory texts it claims to describe.

sentence performs the whole trick in one clause. It converts the child into a "sophisticated bad actor," converts her foreseeable evasion into a liability-severing wrong, and converts a method certification into an outcome shield.

This Article's thesis is that every link in that chain of deflection fails on the deflecting party's own legal materials, and that the question "who is responsible when a child bypasses age verification and is harmed" has a determinate answer that the current commentary, on three continents, has not assembled. The argument runs in three movements.

First, the child cannot be the culpable party, and the reasons are structural rather than sentimental. The age attestation the child supposedly falsified lives inside a clickwrap instrument that never formed as a contract between the platform and the child, because a platform charged with knowledge that its audience includes children cannot, under the objective theory of assent, read a child's tap as a manifestation of agreement to anything; an attestation inside an unformed instrument binds no one and can be breached by no one.[3] Fraud fares no better than contract, because fraud requires justifiable reliance, and every age verification statute on the books exists precisely because the law refuses to credit a child's own statement of age; a regime founded on the premise that a child's word about her age is legally worthless cannot charge the child with fraud for proving the premise correct. And a century of tort doctrine, built in the child labor and dangerous-articles cases and carried into the Restatement, holds that where a statute protects a class against its members' own immaturity, the protected class member's evasion, carelessness, and even misrepresentation of age cannot be turned against her, because the statute was enacted against exactly that conduct.[4]

Second, the bypass cannot be a superseding cause. The very-hazard rule of Restatement (Second) of Torts section 449 provides that where the likelihood of a third person's act is the hazard that makes the actor's conduct negligent, the occurrence of that act does not cut off

3. *See infra* Part IV.A. The non-formation analysis is developed at article length in a companion work; this Article states it as premise and carries it only as far as the age attestation requires.

4. *See infra* Part IV.C.

liability.[5] Children attempting to get past age gates is not an intervening surprise in this system. It is the founding assumption of the system, the reason the gates exist, the behavior the regulators' own guidance instructs platforms to stress-test against. An event cannot simultaneously be the purpose of a duty and the excuse for breaching it.

Third, the compliance shield dissolves in every one of the three jurisdictions this Article surveys, and it dissolves for a different reason in each, which is the comparative finding. In the United States, statutory compliance has never been more than a floor; a defendant who knew of a risk the statute's minimum does not address is negligent for failing to take reasonable measures against it, and the age verification statutes themselves, by attaching enhanced penalties to the event of a minor actually getting through, announce that the legislatures did not regard deployment of the mandated method as the end of the operator's obligations.[6] In the United Kingdom and the European Union, the shield fails on the regimes' own definitions: Ofcom's "highly effective" standard and the European Commission's Article 28 guidelines both build resistance to circumvention into the meaning of effectiveness, with Ofcom expressly instructing services to mitigate "methods of circumvention that are easily accessible to children,"[7] which means a system defeated by makeup, a borrowed account, or a five dollar purchase has not earned the standard whose protection the operator invokes. The comparative punchline inverts the intuition a reader might bring to the survey. The jurisdiction with the least comprehensive child-safety regime, the United States, preserves the least forgiving compliance defense; the jurisdictions with the most comprehensive regimes wrote standards that, correctly read, refuse to attach their protection to any gate that children in fact defeat. The certified-method safe harbor that the industry describes exists in vendor marketing. It does not exist in the law.

The Article proceeds in seven further Parts. Part II surveys the mandate wave across the three jurisdictions and assembles the bypass record. Part III maps the architecture of deflection: the four exits, the scripts each actor reads, and the load-bearing role of the phrase "sophisticated

5. RESTATEMENT (SECOND) OF TORTS § 449 (Am. L. Inst. 1965); *see infra* Part V.

6. *See infra* Part VI.A.

7. *See infra* Parts VI.B–C.

bad actor." Part IV establishes that the child cannot be the culpable party, on contract, fraud, and protective-statute grounds. Part V establishes that the child's bypass is the targeted hazard rather than a superseding cause. Part VI reads the compliance shield against the statutory and regulatory texts of all three jurisdictions and finds it absent from each. Part VII traces where responsibility lands when the deflections fail: on the platform, on the age assurance vendor, and, in the register of policy rather than liability, on the legislative choice to mandate gate rituals rather than safer rooms. Part VIII answers objections.

II. The Mandate Wave and the Bypass Record

A. The United States

American age verification law arrives in two overlapping waves. The first targets sexually explicit content. Texas House Bill 1181, codified at Chapter 129B of the Civil Practice and Remedies Code, requires commercial entities publishing substantial sexual material harmful to minors to "use reasonable age verification methods" to confirm that a visitor is eighteen or older, by digital identification, government-issued identification, or transactional data.[8] The Supreme Court upheld the statute in June 2025, holding that age verification requirements for content obscene to minors trigger and survive intermediate scrutiny, and that the state's "traditional power to prevent minors from accessing speech that is obscene from their perspective . . . includes the power to require proof of age."[9] By the end of 2025, roughly half the states imposed age verification in some form on access to adult content, social media, or both, and several created private rights of action allowing parents to sue noncompliant platforms directly.[10]

8. Tex. Civ. Prac. & Rem. Code §§ 129B.002–.003, https://statutes.capitol.texas.gov/Docs/CP/htm/CP.129B.htm (last visited July 10, 2026).

9. Free Speech Coalition, Inc. v. Paxton, 606 U.S. 461 (2025) (holding that H.B. 1181 "triggers, and survives, review under intermediate scrutiny because it only incidentally burdens the protected speech of adults," and describing government-issued identification and transactional data as "established methods of verifying age").

10. *See* Tex. Civ. Prac. & Rem. Code § 129B.006(b) (2025) (authorizing attorney general civil penalties of up to $10,000 per day of operation in violation of the age verification requirements and, "if, because of the entity's violation of the age verification requirements of this chapter, one or more minors accesses sexual material harmful to minors, an additional amount of not more than $250,000"), https://statutes.capitol.texas.gov/Docs/CP/htm/CP.129B.htm (last visited July

The second wave targets account creation and app distribution. Texas's App Store Accountability Act requires app stores to verify the age of every user at account creation, affiliate each minor's account with a verified parent account, and obtain per-transaction parental consent, and it has been in effect since June 2026, with the Supreme Court declining in July 2026 to disturb the Fifth Circuit's stay of the injunction against it.[11] Beneath both waves sits the federal children's privacy layer: the Children's Online Privacy Protection Act has, since 2000, conditioned collection of personal information from children under thirteen on verifiable parental consent, a regime whose 2025 amendments broadened the definition of personal information to reach biometric identifiers and tightened the consent machinery further.[12]

B. The United Kingdom

The Online Safety Act 2023 requires services that publish or allow pornographic content, and user-to-user services likely to be accessed by children that allow primary priority content, to deploy age assurance that is "highly effective at correctly determining whether or not a particular

---

10, 2026); *Age Verification Bill: Requirements, Risks, and Penalties*, LegalClarity (Apr. 8, 2026), https://legalclarity.org/age-verification-bill-requirements-and-legal-challenges/ (last visited July 10, 2026) (surveying state enactments and Arizona's parallel $250,000 penalty). Private enforcement varies by state. Louisiana, the first mover, permits an individual to sue a commercial entity where the verification failure "results in a minor accessing harmful material"; Virginia makes noncompliant entities liable for the damages resulting from a minor's access, plus attorney fees, and Virginia and Utah provide for private enforcement only; Arkansas, Mississippi, Montana, and North Carolina likewise expose noncompliant websites to damages actions by individuals. Texas, by contrast, commits enforcement of its age verification requirement solely to the attorney general, reserving its private action for unlawful retention of identifying information. *See* Eric N. Holmes, Cong. Rsch. Serv., LSB11020, *Online Age Verification (Part I): Current Context* (2023), https://www.congress.gov/crs_external_products/LSB/HTML/LSB11020.web.html (last visited July 10, 2026) (Texas's attorney general "has sole enforcement authority" over the state's age verification law, while "Virginia's and Utah's pornography age verification laws provide only for enforcement by individuals"); La. Rev. Stat. § 9:2800.29(B)(3)(a) (enacted by 2022 La. Acts 440, eff. Jan. 1, 2023) (a violating entity "shall be liable to an individual for damages resulting from a minor's accessing the material, including court costs and reasonable attorney fees as ordered by the court"), https://legis.la.gov/legis/Law.aspx?d=1293341 (last visited July 10, 2026); Va. Code § 8.01-40.5; *Age Verification Laws by State: What They Require*, LegalClarity (May 28, 2026), https://legalclarity.org/age-verification-laws-by-state-what-they-require/ (last visited July 10, 2026).

11. Tex. Bus. & Com. Code §§ 121.021–.022; Students Engaged in Advancing Tex. v. Paxton, No. 25A1389 (U.S. July 6, 2026); Comput. & Commc'ns Indus. Ass'n v. Paxton, No. 25A1390 (U.S. July 6, 2026) (orders denying applications to vacate the stay).

12. 15 U.S.C. §§ 6501–6506; 16 C.F.R. pt. 312; Children's Online Privacy Protection Rule, 90 Fed. Reg. 16918 (Apr. 22, 2025).

user is a child."[13] Ofcom's implementing guidance, published in January 2025, gives the phrase operational content through four criteria: an age assurance process must be technically accurate, robust, reliable, and fair.[14] Two of the four criteria matter enormously for this Article and are quoted precisely in Part VI: technical accuracy measures performance under laboratory conditions, while robustness measures performance "in actual deployment contexts," including resistance to circumvention, with Ofcom instructing that a photo-identification method "should not be capable of being easily circumvented" and should detect "falsified documentation or manipulation that a child could create or obtain."[15] Self-declaration is ruled out as incapable of being highly effective. Compliance deadlines have passed; enforcement is running; fines reach the greater of eighteen million pounds or ten percent of global revenue.

C. The European Union

Article 28(1) of the Digital Services Act obliges online platforms accessible to minors to "put in place appropriate and proportionate measures to ensure a high level of privacy, safety, and security of minors."[16] The European Commission's guidelines under Article 28(4), published July 14, 2025, recommend age assurance methods that are accurate, reliable, robust, non-intrusive, and non-discriminatory, prescribe age verification for adult content and for services subject to national minimum-age rules, and state flatly that self-declaration does not meet the

13. Online Safety Act 2023, c. 50 (UK); *see* Ofcom & Info. Comm'r's Off., *Joint Statement on Age Assurance* (Mar. 2026), *summarized in Ofcom and ICO Issue Joint Statement on Age Assurance*, Covington, Inside Privacy (Apr. 1, 2026), https://www.insideprivacy.com/online-safety/ofcom-and-ico-issue-joint-statement-on-age-assurance/ (last visited July 10, 2026).

14. Online Safety Act 2023, c. 50, § 81(3) (UK); Ofcom, *Guidance on Highly Effective Age Assurance and Other Part 5 Duties* ¶¶ 4.1, 4.23–4.24 (Jan. 16, 2025), https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10- (last visited July 10, 2026) ("the age verification or age estimation must be of such a kind, and used in such a way, that it is highly effective at correctly determining whether or not a particular user is a child").

15. Ofcom, *supra* note 14, ¶¶ 4.29, 4.44, 4.57; *see also UK Online Safety Act: Ofcom Publishes Guidance on Age Assurance and Children's Access Assessments*, Osborne Clarke (Feb. 4, 2025), https://www.osborneclarke.com/insights/uk-online-safety-act-ofcom-publishes-guidance-age-assurance-and-childrens-access (last visited July 10, 2026).

16. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 (Digital Services Act), art. 28(1).

standard.[17] Spain's coordinating authorities gloss the robustness criterion with a question that could serve as this Article's epigraph: "How easy is it for a tech-savvy teen to circumvent the age restriction?"[18] The guidelines carry a feature that platforms cite less often than they should: they are not legally binding, and following them "does not automatically guarantee compliance" with Article 28(1).[19] Australia completes the survey's outer edge with an outright statutory ban on social media accounts for children under sixteen, in force since 2025.[20]

D. The Bypass Record

Against this legal architecture, set the empirical one. In January 2026, after litigation pressure and state investigations, Roblox mandated facial age verification worldwide for chat access, sorting users into six age brackets.[21] Within days, age-verified accounts appeared for sale on eBay for roughly five dollars, allowing any purchaser, including the adult predators the brackets exist to wall out, to enter child-designated spaces carrying a pre-validated minor status.[22] In the same month, users reported inaccurate age estimates from the biometric system and documented ways around it.[23] In Australia, children continue to get around the under-sixteen ban in numbers, a fact reported matter-of-factly in coverage of the regime's first year.[24] And beneath the headline systems sits the ordinary infrastructure: empirical testing of the platforms most popular with

---

17. European Comm'n, *Guidelines on Measures to Ensure a High Level of Privacy, Safety and Security for Minors Online* (July 14, 2025), https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-protection-minors (last visited July 10, 2026); *see The Long-Awaited EU Guidelines on Article 28(1) DSA*, HLC (Aug. 5, 2025), https://www.hlc.com/en/publications/the-longawaited-eu-guidelines-on-article-281-dsa-what-online-platforms-must-know (last visited July 10, 2026) (the method must be "robust enough to prevent circumvention").

18. CNMC & AEPD, *Decoding Article 28 of the DSA: Age Assurance and Service Design for Online Platforms* (Mar. 24, 2026), https://www.aepd.es/en/press-and-communication/blog/decoding-article-28-dsa-age-assurance-service-design-online-p (last visited July 10, 2026).

19. European Comm'n, *supra* note 17 ("following these guidelines is voluntary and does not automatically guarantee compliance").

20. Online Safety Amendment (Social Media Minimum Age) Act 2024 (Cth) (Austl.).

21. Anapol Weiss, *supra* note 1.

22. *Id.*

23. Angela Yang, *Age Verification Is Coming for the Internet*, NBC News (Apr. 16, 2026), https://www.nbcnews.com/tech/tech-news/age-verification-laws-advocates-express-concerns-rcna331835 (last visited July 10, 2026) ("After Roblox began requiring biometric age checks to access chat, some users reported inaccurate age estimates, while others found ways to bypass the system.").

24. *Id.* ("in Australia, many continue to get around the country's social media ban for users under 16").

European minors, published in October 2025, found that every one of the eight services tested, including Discord, Fortnite, Instagram, Roblox, Snapchat, TikTok, Twitch, and YouTube, relied at account creation on self-declaration, the one method every regulator on every continent has ruled insufficient, and that where parental consent tools existed they could be trivially evaded.[25]

Three features of this record control the legal analysis that follows, so they deserve plain statement. First, the bypasses are not exotic. A birthdate field accepts arithmetic; a facial estimator near its published error band accepts makeup and older siblings; a bracket system accepts a purchased account. These are methods available to children, executed by children, at scale. Second, the bypasses are documented within days of each system's deployment, which means the operators possess actual knowledge of the failure modes for essentially the entire operating life of each gate. Third, and consequently, the population coming through the gate is the population the gate was built for. Nothing about the record resembles the "highly sophisticated bad actor" of the industry's liability narrative. The record shows the gate's own threat model walking through the gate.

III. The Architecture of Deflection

When the harm arrives, four actors read four scripts, and the scripts interlock so that each actor's exit depends on the next actor's presence. The structure deserves to be mapped before it is dismantled, because its power lies in the interlocking rather than in any single argument.

The platform's script is compliance. It deployed the method the statute or the regulator's guidance named; it engaged a certified vendor; the gate was there and the gate was blessed. Whatever happened past the gate happened despite the platform's discharge of the only duty the mandate imposed. The script's legal form is the statutory-compliance defense in the United States and the "highly effective age assurance" designation in the United Kingdom and the European Union, and its commercial form is the safe harbor the age assurance industry markets to its

25. *Mind the Gap: Age Assurance Under the DSA* (interface, Oct. 13, 2025), https://www.interface-eu.org/publications/age-assurance-g (last visited July 10, 2026).

customers: deploy a certified system, and "if a highly sophisticated bad actor bypasses" it, "the platform is generally protected from liability."[26]

The vendor's script is fraud. Its system performed within specification; the user defeated it by deception, presenting a face, a document, or an account that was not what it purported to be. Deception by the verified party is the responsibility of the deceiver. The script's legal form is the superseding-cause doctrine, with the bypasser cast as the intervening wrongdoer whose act severs the chain, and its contractual form is the indemnification and limitation clauses that quietly route any residual vendor exposure back to the platform.

The state's script is silence. The legislature that mandated the gate, selected the method categories, and blessed the certification scheme cannot be sued for the gate's design. Legislative immunity is not wrong as law, and this Article does not contest it. Its role in the architecture is structural: the state's unreachability is what makes the platform's "we did what the law required" script terminal. If the law required a defective gate, and the law's author cannot answer, the defect has no owner.

The child's script is written for her by the other three. She is the residual. The platform's compliance, the vendor's fraud theory, and the state's immunity leave exactly one actor at the scene, and the vocabulary is already prepared: she *misrepresented* her age; she *circumvented* a security system; she *deceived* the verification. The eleven year old in her sister's makeup becomes, in the vendor's white paper, a sophisticated bad actor, and in the platform's motion, an intervening wrongdoer, and the harm she suffered inside the space she was never supposed to reach becomes, at the end of the chain of scripts, her own fault.

Notice what the entire architecture requires. Every script depends on one characterization: that the child's passage through the gate was a legally cognizable wrong, a fraud, a deception, an act culpable enough to absorb the responsibility the other actors are shedding. The platform's compliance is only terminal if the bypass was someone else's wrong. The vendor's specification

26. Klippa, *supra* note 2.

defense only works if defeat of the specification was wrongful rather than foreseeable. The word "fraud" is load-bearing for the whole structure. Parts IV and V remove it.

IV. The Child at the Gate

A. Nothing Was Formed: The Attestation Inside an Instrument That Does Not Exist

Begin with the object the child supposedly falsified. When a platform accuses a child of misrepresenting her age, the representation it points to almost always lives inside its own onboarding flow: the birthdate field, the "I am over 13" checkbox, the terms-of-service acceptance that incorporates an age warranty. The accusation therefore presupposes an instrument, a formed agreement between the platform and the child within which the age term has legal life. That presupposition fails, and it fails on the platforms' own preferred doctrine.

Mutual assent is measured objectively, from the position of the offeror, charged with what it knows and has reason to know about the party manifesting assent. The objective theory's canonical statement carries its own limiting principle: a party's undisclosed intention is immaterial "except when an unreasonable meaning which he attaches to his manifestations is known to the other party."[27] Objectivity protects an offeror who reasonably relies on an outward manifestation; it withdraws that protection from an offeror who knows or has reason to know that the manifestation does not carry the meaning the offeror assigns it. A platform that designs child-attractive products, markets into child audiences, monetizes child engagement, and operates under a federal statute that has obliged it since 2000 to know whether its service is directed to children[28] cannot reasonably read a child's tap as a manifestation of assent to thousands of words of undisclosed terms. The inquiry-notice standard the online contract cases apply is calibrated to a "reasonably prudent" user whose every competence, recognizing hyperlink conventions, inferring juridical meaning from spatial layout, understanding that a button labeled Continue

27. Lucy v. Zehmer, 196 Va. 493, 503 (1954).
28. 15 U.S.C. §§ 6501–6506; 16 C.F.R. pt. 312.

performs a second and invisible function of legal commitment, is an adult competence,[29] and it returns no answer at all when the user is nine. The First Circuit has already held that an interface that fails to communicate the existence of terms to its actual audience forms no agreement, because "[t]here can be no mutual assent or meeting of the minds . . . if the one to whom the offer is supposedly made is unaware of the contract's terms and conditions,"[30] and that a defect of this kind is a formation defect, reserved to the court and beyond the reach of any delegation clause.[31] This Article develops that argument no further, because a companion work develops it at length; here it operates as premise.

The premise's consequence for the fraud narrative is total. An age attestation inside an instrument that never formed is not a warranty, not a contractual representation, not a term the child breached. It is a string the child typed into a box, on a screen that communicated "game" to its known audience, at the invitation of an operator who designed the box to accept any input and verify nothing. There is no agreement containing the representation; there is nothing for the representation to be *inside of*. The platform that says "she agreed she was eighteen" is pointing at a contract that does not exist to prove a wrong that therefore cannot exist either.

B. The Reliance the Law Forbids

Suppose the platform retreats from contract to tort and pleads the child's misstatement of age as common law fraud or as the basis for estoppel. Fraud requires, among its elements, justifiable reliance by the party claiming deception. Here the entire statutory environment speaks with one voice, and what it says is fatal: the operator is not permitted to rely on a child's own statement of her age, and it never was.

---

29. *See* Meyer v. Uber Techs., Inc., 868 F.3d 66, 74–80 (2d Cir. 2017) (calibrating inquiry notice to what "a reasonably prudent smartphone user" knows about hyperlinks and interface conventions).

30. Nat'l Fed'n of the Blind v. Container Store, Inc., 904 F.3d 70, 82–83 (1st Cir. 2018).

31. *Id.*; Buckeye Check Cashing, Inc. v. Cardegna, 546 U.S. 440, 444 n.1 (2006) (distinguishing validity from "whether any agreement between the alleged obligor and obligee was ever concluded"); Granite Rock Co. v. Int'l Bhd. of Teamsters, 561 U.S. 287, 297 (2010); *cf.* K.F.C. v. Snap Inc., 29 F.4th 835, 837 (7th Cir. 2022) ("Even the most sweeping delegation cannot send the contract-formation issue to the arbitrator, because, until the court rules that a contract exists, there is simply no agreement to arbitrate.").

Every enactment surveyed in Part II is, at bottom, a legislative finding that a child's self-attestation of age is legally worthless. COPPA has conditioned data collection on *verifiable parental* consent for a quarter century precisely because the child's own click settles nothing.[32] Ofcom rules self-declaration out of the class of methods capable of being highly effective.[33] The European Commission "explicitly states that self-declaration does not meet these standards."[34] Texas requires verification by identification or transactional data for exactly the reason that a birthdate field is not evidence of anything.[35] A regime whose founding premise is that the child's word about her age carries no legal weight cannot furnish the justifiable reliance that fraud requires when the child's word turns out, as the premise predicted, to be unreliable. The operator did not justifiably rely; the operator was forbidden to rely, on pain of the very statutes it now invokes as its shield. The syllogism can be stated in a sentence: no one may found a legal regime on the worthlessness of a child's self-attestation and then prosecute the child's self-attestation as fraud. The regime and the fraud theory cannot both be true, and the regime is on the books.

The point generalizes past the birthdate field to the biometric and account-based bypasses. The eleven year old who presents her made-up face to an estimator is not making a representation the law recognizes at all; she is presenting her face, and the estimator's judgment about that face is the vendor's output, produced by the vendor's model, within the vendor's published error bands. A system whose specification sheet concedes a margin of error at the threshold has priced in the made-up eleven year old; her arrival within the error band is the specification performing as documented, and a party cannot be defrauded by the occurrence of its own specification. The purchased-account bypass differs, and Part VII returns to it, because the purchaser is an adult with genuine scienter. But even there, the fraud runs against the verification system while the harm lands on a child, and the doctrine has a name for defendants who invoke a third party's wrong to escape a duty that existed because of that wrong. Part V supplies it.

---

32. 15 U.S.C. § 6501–6506; 16 C.F.R. § 312.5.
33. Ofcom, *supra* note 14.
34. HLC, *supra* note 17.
35. Tex. Civ. Prac. & Rem. Code § 129B.003.

C. The Protected Class Cannot Be Charged With Defeating Its Protection

American tort law confronted this exact configuration a century ago, in factories and gun shops rather than platforms, and resolved it with a doctrine the platform cases have not yet been forced to meet. Where a statute is enacted to protect a class of persons against their own inability to protect themselves, against their own "inexperience, lack of judgment, and tendency toward negligence," the courts construe the statute to place the entire responsibility on the regulated party, and the protected class member's own carelessness cannot be turned against her.[36] The paradigm statutes are the child labor acts and the laws prohibiting the sale of firearms and other dangerous articles to minors. The New York Court of Appeals stated the rule's logic in terms that transfer to the age gate without modification: recovery by a child protected by such a statute "shall not be defeated by the very negligence, lack of care and caution that the statute was designed to prevent and make impossible."[37] The child's imprudence is not a defense to the violation of a statute that exists because children are imprudent. To hold otherwise "would be to defeat the evident purpose of the statute."

Two features of this line of cases map onto the age verification setting with uncomfortable precision for the platforms. The first is that the doctrine survives the minor's own misrepresentation of age. In *Tamiami*, the Florida Supreme Court held a gun dealer strictly responsible for injury to a sixteen year old who had bought a rifle representing himself as "over 18," explaining that such statutes make the seller "liable for injury to the child even though he has acted in good faith and has employed the infant in ignorance of his age."[38] The second is that the regulated party carries an affirmative duty of vigilance about age that a bare representation

---

36. RESTATEMENT (SECOND) OF TORTS § 483 cmt. c (Am. L. Inst. 1965); Tamiami Gun Shop v. Klein, 116 So. 2d 421, 423 (Fla. 1959) (statutes protecting a class "from their inability to exercise self-protective care" are "construed to place the entire responsibility upon the defendant, and to require him to protect not only plaintiffs who are exercising reasonable care but those who are contributorily negligent as well").

37. Vincent v. Riggi & Sons, Inc., 30 N.Y.2d 406 (1972) (quoting Karpeles v. Heine, 227 N.Y. 74, 80 (1919)) (holding that violation of a child labor statute establishes fault regardless of the child's contributory negligence).

38. *Tamiami*, 116 So. 2d at 423 (adopting the child labor analogy); *see also* Del E. Webb Corp. v. Superior Court, 726 P.2d 580 (Ariz. 1986) (describing *Tamiami*: the statute "precluded the defenses of contributory negligence and assumption of the risk" notwithstanding the minor's representation that he was over eighteen).

cannot discharge. The New York rule, quoted at length in *Vincent*, holds that "[t]he representation of the employee as to his age . . . is not conclusive on the question. No principle of estoppel is applicable to the case. The question always is whether the employer is justified in believing that the employee is of sufficient age," and the employer "may not rest alone on the representation of the plaintiff, but is required to exercise proper vigilance to discover the fact."[39] A doctrine written in 1909 for candy factories states the modern age verification duty more honestly than the modern statutes do: the child's own answer to the age question is never the end of the regulated party's obligation, no estoppel arises from it, and vigilance proportionate to the circumstances is the standard. For the youngest children the law goes further still and forecloses the inquiry entirely; a very young child is conclusively presumed incapable of negligence at all.[40]

Assemble the three subsections and the child's position at the gate is legally determinate. Contract law finds no instrument within which her attestation could be a breach. Fraud law finds no reliance the operator was permitted to place. Tort law's protective-statute doctrine affirmatively bars charging her evasion against her, on the strength of a century of cases holding that the class a statute protects from its own immaturity cannot be met with that immaturity as a defense. The eleven year old in her sister's makeup is not a sophisticated bad actor. She is the statutory subject, doing the statutorily anticipated thing, and every body of law that touches her says the same sentence: the gate was the grown-ups' job.

V. The Bypass Is the Hazard, and the Hazard Cannot Supersede

With the child's culpability removed, the deflection architecture needs the bypass itself, however characterized, to sever causation: the platform and vendor argue that whatever their antecedent failures, the act of getting through the gate was an intervening event that breaks the chain between their conduct and the harm. Tort law rejected this move, in general form, before any of these companies existed.

---

39. *Vincent*, 30 N.Y.2d 406 (quoting Koester v. Rochester Candy Works, 194 N.Y. 92 (1909)).

40. Miller v. Warren, 390 S.E.2d 207 (W. Va. 1990) (syllabus) ("A two-year-old child is conclusively presumed incapable of negligence."); *see* RESTATEMENT (SECOND) OF TORTS § 283A (Am. L. Inst. 1965).

The very-hazard rule, stated at Restatement (Second) of Torts section 449 and adopted across the states, provides: "If the likelihood that a third person may act in a particular manner is the hazard or one of the hazards which makes the actor negligent, such an act whether innocent, negligent, intentionally tortious, or criminal does not prevent the actor from being liable for harm caused thereby."[41] Its companion, section 302B, makes the point from the duty side: an act or omission may be negligent precisely because the actor should realize it creates an unreasonable risk of harm through the conduct of a third person, "even though such conduct is criminal."[42] The rule's logic is a conservation law for responsibility: an event cannot be simultaneously the reason a precaution is required and the excuse for the precaution's failure. The landlord who leaves the trash piled against the wall does not escape when the fire is set by an arsonist, because fire, however ignited, is the hazard the trash rule addresses.[43]

Now ask the section 449 question about the age gate, and notice that it barely needs to be argued, because the regulators have already answered it in writing. What is the hazard that makes a defective age verification system negligent to deploy? Children attempting to pass it. That is the entire and only hazard. The system has no other purpose; the duty has no other content. Ofcom's robustness criterion instructs services to stress-test "against common circumvention attempts," to detect "falsified documentation or manipulation that a child could create or obtain," and to "identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children."[44] The European guidance defines robustness by asking how easily a teenager can get around the system.[45] The duty is, in the regulators' own words, a duty

41. RESTATEMENT (SECOND) OF TORTS § 449 (Am. L. Inst. 1965); Britton v. Wooten, 817 S.W.2d 443 (Ky. 1991) (adopting §§ 448–449 and § 302B and observing that the Restatement appendices collect "perhaps one hundred" cases so holding; an intervening act, even criminal, supersedes only where it is so "highly extraordinary" that the antecedent negligence should be ruled out as a substantial factor).

42. RESTATEMENT (SECOND) OF TORTS § 302B (Am. L. Inst. 1965); *Britton*, 817 S.W.2d 443.

43. *Britton*, 817 S.W.2d 443 (whether the spark "was ignited negligently, intentionally, or even criminally" is not the issue where the accumulation negligently contributed to the spread).

44. Ofcom, *supra* note 14, ¶¶ 4.44–4.45, 4.52, 4.57; Ofcom, *Highly Effective Age Assurance: Frequently Asked Questions*, https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/enforcement/highly-effective-age-assurance–frequer (last visited July 10, 2026).

45. CNMC & AEPD, *supra* note 18.

calibrated to the child's anticipated evasion. Under section 449, that anticipated evasion cannot be a superseding cause of the harm that follows it. The child getting through is not the chain breaking. The child getting through is the chain.

The same analysis absorbs the adult who buys the five dollar account, the one genuine wrongdoer in the record, and this is where the deflection architecture's strongest piece fails. His fraud is real; his scienter is real; nothing in Part IV shelters him, and he should answer for his own conduct in every forum that can reach him. But his intervening crime does not launder the platform's antecedent negligence, for two independent reasons. First, predators seeking entry to child spaces are, alongside children seeking exit from child restrictions, the other named hazard of the bracket system; a verification architecture whose selling point is walling adults out of child rooms cannot treat an adult's entry into a child room as an unforeseeable bolt from outside the risk. Second, and dispositively on these facts, the secondary market in verified accounts was documented within days of the system's launch and persisted in public view.[46] From that point forward the operator possessed specific knowledge, well above any baseline of generalized risk, that its gate was being sold at retail, and it continued operating the gate as its compliance showpiece. An operator with specific knowledge of an active failure mode who continues the conduct occupies the least sympathetic position known to negligence law; the intervening criminal act it now pleads as superseding is one it watched being advertised. Section 449 was written for less.

VI. The Compliance Shield, Read Against Its Own Text

The last standing deflection is the platform's: whatever tort law might otherwise say, the operator did what the mandate required, and compliance discharges the duty. This Part reads that defense against the governing texts of each jurisdiction and finds that it fails in all three, for three different reasons, and that the differences are the comparative finding this Article set out to establish.

---

46. Anapol Weiss, *supra* note 1.

A. The United States: The Floor That Never Became a Ceiling

The American rule is old, general, and stated in the Restatement: a statutory or regulatory standard of conduct "is normally a minimum standard," and the "legislative or administrative minimum does not prevent a finding that a reasonable [person] would have taken additional precautions where the situation is such as to call for them."[47] Its sharpest formulation comes from West Virginia, in language that could have been drafted for the five dollar account market: "Compliance with a regulation does not constitute due care per se," and "[i]f the defendants knew or should have known of some risk that would be prevented by reasonable measures not required by the regulation, they were negligent if they did not take such measures."[48] A platform that deployed the mandated verification method and then learned, within days, that verified child-status accounts were selling for five dollars, that its estimator was waving through made-up eleven year olds, and that its brackets were being crossed in both directions, holds exactly the knowledge *Miller* describes: knowledge of a risk that reasonable measures beyond the mandate, invalidating resold accounts, re-verifying on device changes, rate-limiting bracket privileges, hardening the estimator's challenge age, would prevent. Compliance with the statute is evidence for the jury. It is not an answer.

The honest treatment of this rule requires acknowledging its exception, because the platforms will plead it. Courts will occasionally accept statutory compliance as fixing the standard of care where the legislature or regulator has specifically deliberated the precise question and the case presents only "the ordinary situation contemplated" by the enactment.[49] The exception does not reach these facts, and the age verification statutes say so themselves, in their penalty structures. Texas attaches an additional civil penalty of up to $250,000 to the event

---

47. RESTATEMENT (SECOND) OF TORTS § 288C (Am. L. Inst. 1965), *quoted in* Ramirez v. Plough, Inc., 6 Cal. 4th 539, 548 (1993) ("Courts have generally not looked with favor upon the use of statutory compliance as a defense to tort liability.").

48. Miller v. Warren, 390 S.E.2d 207 (W. Va. 1990) (syllabus and accompanying text; "a statute or regulation merely sets a floor of due care").

49. *Ramirez*, 6 Cal. 4th at 548 (adopting the regulatory standard as the tort standard for the specific, expressly deliberated question of foreign-language drug labeling).

of one or more minors actually accessing the restricted material because of the violation, over and above the daily penalties for the violation itself, and Arizona replicates the structure while going one step further: it hands the enhanced penalty to the parent of the child who got through, as a private right of action with attorney fees.[50] Arizona is the sharpest drafting in a wider pattern: Louisiana, the first mover, lets an individual sue when the verification failure results in a minor accessing the material; Virginia makes the noncompliant entity liable for the damages resulting from the minor's access, with attorney fees; Kentucky awards $10,000 per instance that the platform failed to verify, recoverable by the injured person or by a parent or legal guardian suing on the injured minor's behalf, plus actual damages and fees; and Virginia and Utah provide for private enforcement only.[51] A legislature that prices the child-gets-through event separately from the no-gate-deployed event has told the courts, in operative text, that deploying the gate was never understood to end the operator's obligations; the statutes themselves contemplate, and monetize, the mandated method's failure. There is no *Ramirez*-style settled deliberation to defer to, because the enactments' own structure treats bypass as a live, priced, expected residual risk. The floor, in other words, was built with a trapdoor in it, and the legislatures labeled the trapdoor.

A caveat about the mandates' constitutional fates belongs here, because it changes less than it appears to. The two American waves have fared differently in court. Verification tied to material obscene as to minors survived *Paxton* under intermediate scrutiny; the social media account statutes have been repeatedly enjoined as content based, with Louisiana's Secure Online Child Interaction and Age Limitation Act falling most recently to a permanent injunction in

50. Tex. Civ. Prac. & Rem. Code § 129B.006(b)(1), (3) (2025); Ariz. Rev. Stat. § 18-701(F), (H) ("[t]he parent or guardian of a minor who accesses material harmful to minors . . . has a right of action against the offending entity," with a penalty of up to $10,000 per day of violation and "an additional amount of not more than $250,000 if, because of the entity's violation of the age verification requirements of this section, one or more minors accesses sexual material that is harmful to minors," plus reasonable attorney fees and costs); LegalClarity, *supra* note 10 (surveying the other private-action states).

51. *Supra* note 10; Ky. Rev. Stat. § 436.002 (created by 2024 Ky. Acts ch. 106, § 14, eff. July 15, 2024) ("[a]ny person injured by a violation of this section, or a parent or legal guardian on behalf of any minor injured by a violation of this section, may bring a civil action against the covered platform to recover . . . [d]amages of ten thousand dollars ($10,000) per instance that the covered platform failed to perform age verification to restrict the minor's access to matter harmful to minors" as well as "[a]ctual damages, court costs, and reasonable attorney's fees").

December 2025, an injunction that leaves the 2022 Louisiana adult-content statute cited above untouched.[52] A struck mandate removes a negligence per se predicate. It does not dissolve the common law duty, because that duty attaches to any gate an operator in fact deploys, mandated or voluntary, and the deployed gates of the bypass record, Roblox's first among them, were adopted under litigation pressure rather than statutory command. The duty follows the deployment, and the deployments are not receding.

One American statute does contain a genuine method-based defense, and precision about it strengthens the account rather than weakening it. The Texas App Store Accountability Act provides that app stores are not liable for breaches of its age verification and parental consent requirements where they use widely adopted industry standards and apply those standards consistently and in good faith; Utah's twin statute, enacted the same season on the same architecture, contains no safe harbor at all, and instead hands the minor or parent a private action for the greater of actual damages or $1,000 per violation, with attorney fees.[53] Three boundaries keep the Texas provision from growing into the industry's imagined shield. It is internal to that Act, a defense against that Act's own obligations rather than against common law negligence for independently known hazards. It is conditioned on good faith and consistency, and those conditions do not survive an operator's specific knowledge of an operating failure mode; good faith is a description of ignorance, and the bypass record of Part II is a record of knowledge.

---

52. NetChoice, LLC v. Murrill, No. 3:25-cv-00231 (M.D. La. Dec. 15, 2025) (granting summary judgment and permanently enjoining enforcement of Louisiana's Secure Online Child Interaction and Age Limitation Act, 2023 La. Acts 456 (S.B. 162), against NetChoice's members); *see also* NetChoice, LLC v. Griffin, No. 5:23-cv-05105 (W.D. Ark. Mar. 31, 2025) (permanently enjoining Arkansas's Social Media Safety Act); NetChoice, LLC v. Yost, No. 2:24-cv-00047 (S.D. Ohio Apr. 16, 2025) (same as to Ohio's Parental Notification by Social Media Operators Act); *Louisiana Social Media Age Gate Law Hit by Permanent Injunction*, Bloomberg Law (Dec. 16, 2025), https://news.bloomberglaw.com/us-law-week/louisiana-social-media-age-gate-law-hit-by-permanent-injunction (last visited July 10, 2026).

53. Tex. Bus. & Com. Code ch. 121; Utah Code Ann. §§ 13-76-201 to -202 (2025), https://le.utah.gov/xcode/Title13/13.html (last visited July 10, 2026); *see Texas and Louisiana Join Growing Trend of State Age Verification Laws for App Stores*, Orrick (July 10, 2025), https://www.orrick.com/en/Insights/2025/07/Texas-and-Louisiana-Join-Growing-Trend-of-State-Age-Ve (last visited July 10, 2026) (describing the Texas Act's provision that app stores are not liable "if they use widely adopted industry standards to verify age and obtain parental consent and apply those standards consistently and in good faith," and noting that this "contrasts with the Utah's App Store Accountability Act, which does not provide a safe harbor" and that Utah's law is enforceable "by a minor or parent through a private right of action, for the greater of actual damages or $1,000 for each violation, reasonable attorney fees and costs").

And its omission from the Utah twin shows that legislatures know exactly how to write the shield when they intend one, which sharpens the inference against every statute that lacks it.

One more American peculiarity requires a paragraph, because it is the objection every platform lawyer reaches for first: Section 230 of the Communications Decency Act. It does not reach this claim. The negligence theory developed here faults no third-party content and treats the platform as the publisher of nothing. It faults the platform's own conduct in designing, deploying, and continuing to operate a defective verification gate, and the courts of appeals have now twice drawn exactly the line that lets such a claim through. The Ninth Circuit holds that Section 230 "cuts off liability only when a plaintiff's claim faults the defendant for information provided by third parties," and that a duty "to design a reasonably safe product is fully independent" of the platform's role in publishing content.[54] The Third Circuit holds that platforms are immunized only for third-party speech and answer for their own first-party conduct and expression.[55] An age gate is not speech at all; it is access-control machinery, the platform's own product feature, and a claim that the machinery was negligently built and knowingly left broken is a *Lemmon* claim in its purest form.

B. The United Kingdom: The Standard That Contains Its Own Refutation

The British compliance story the industry tells runs: Ofcom blessed a menu of methods; the operator picked from the menu; the operator is therefore highly effective; bypass is the bypasser's problem. The story fails at its second step, because "highly effective" is not a designation a method carries on its label. It is a performance property the deployed process must actually have, defined by four criteria, and the criterion the bypass record speaks to is robustness: "the degree to which an age assurance method can correctly determine the age of a user in actual deployment contexts."[56] Ofcom's guidance makes the criterion's content explicit. Robustness is threatened

54. Lemmon v. Snap, Inc., 995 F.3d 1085 (9th Cir. 2021) (negligent design claim predicated on the platform's own Speed Filter feature not barred; the platform is sued "for its own conduct").

55. Anderson v. TikTok, Inc., 116 F.4th 180 (3d Cir. 2024) (platforms "are not immunized if they are sued for their own expressive activity or content").

56. Ofcom, *supra* note 14, ¶ 4.44.

by "circumvention techniques that are easily accessible to children," the guidance's illustrative example being "a child user uploading an image of an ID that does not belong to them"; a robust photo-identification check "should not be capable of being easily circumvented" and must detect "falsified documentation or manipulation that a child could create or obtain"; and services are expected "to identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children and where it is reasonable to assume that children may use them."[57] The guidance then forecloses the menu-selection and certification readings by name, twice. Implementing one of the listed methods "is not a guarantee that the service is acting in accordance with the requirements of the Act,"[58] and using a service certified against a standard or scheme "is not an automatic means of compliance."[59] And Ofcom's own published examples of non-compliance describe the bypass record of Part II with uncomfortable precision: facial age estimation "which allows children to upload a still image they have obtained of an adult"; processes with "no way of ensuring that the details provided . . . belong to the user attempting to access the service"; photo-identification matching "which easily allows children to verify their age using fake or manipulated ID documents."[60] Commentary on the enforcement program draws the inference this Article needs drawn: "If age verification or age estimation set ups are easily tricked then arguably this is not highly effective."[61]

Read that way, and there is no other way to read it, the British standard does something remarkable to the compliance defense: it makes the bypass record and the compliance claim mutually exclusive. A system that eleven year olds defeat with makeup, that children defeat with borrowed accounts, that a five dollar purchase defeats categorically, is a system failing the robustness criterion in deployment, whatever its laboratory accuracy scores said. The

---

57. Ofcom, *supra* note 14, ¶¶ 4.44–4.45, 4.52, 4.57; *see also id.* ¶ 4.58 (repeated age checks can "prevent child access via device or account sharing").

58. *Id.* ¶ 4.8.

59. *Id.* ¶ 4.27.

60. *Id.* ¶¶ 4.52–4.60 & accompanying examples of non-compliance.

61. *The Online Safety Act: Children's Duties, Age Verification and Content Moderation*, Online Safety Act.net, https://www.onlinesafetyact.net/analysis/the-online-safety-act-childrens-duties-age-verification-and-content-moderation-on-user-to-user (last visited July 10, 2026).

demonstrated child-level bypass is not an event the highly effective designation excuses. It is evidence the designation was never earned. The operator who waves the "highly effective age assurance" banner over a gate with a documented resale market is not invoking a safe harbor; it is confessing, in the regulator's own vocabulary, to the element of the violation. The safe harbor the vendor marketing describes, protection so long as only "sophisticated" actors get through, exists nowhere in Ofcom's text. What the text contains is the opposite: a standing duty of circumvention vigilance calibrated to what children can obtain, which is to say, a codified version of the century-old *Koester* rule that the regulated party may never rest on the child's presentation and must exercise proper vigilance to discover the fact.[62]

C. The European Union: Guidance That Guarantees Nothing

The European version of the shield fails even earlier, at the threshold. Article 28(1) imposes the outcome duty, "appropriate and proportionate measures to ensure a high level of privacy, safety, and security of minors," and the Commission's guidelines, the only document a platform could plausibly wave as certification of its methods, disclaim the shield function in terms: the guidelines are voluntary, and following them "does not automatically guarantee compliance" with Article 28(1).[63] What the guidelines do contain is the same robustness criterion as the British standard, glossed by national authorities as the question of how easily a teenager circumvents the system, alongside the flat exclusion of self-declaration.[64] The European operator therefore holds the weakest version of the defense: a non-binding reference document, which conditions even its persuasive force on a circumvention-resistance criterion the bypass record contradicts, deployed against an outcome duty the harm itself impeaches. Meanwhile the empirical baseline across the platforms most used by European minors remained, as of late 2025, self-declaration,

62. *See supra* note 39 and accompanying text.

63. European Comm'n, *supra* note 17.

64. CNMC & AEPD, *supra* note 18; HLC, *supra* note 17 ("robust enough to prevent circumvention").

the one method the guidelines rule out by name.[65] For that population of services the compliance conversation never begins.

D. The Comparative Finding

Set the three readings side by side. The United States, with the thinnest and most fragmented child-safety regime of the three, preserves the strongest residual liability: statutory compliance is a floor, the statutes price the bypass event separately, private rights of action exist, and Section 230 does not reach design conduct. The United Kingdom and the European Union, with the most comprehensive regimes, are widely believed, and marketed, to have traded that comprehensiveness for a certified-method safe harbor. The belief is false on the texts. Both regimes define effectiveness partly as resistance to child-accessible circumvention, which means both regimes convert every documented bypass into evidence against the operator rather than absolution for it, and the European guidelines disclaim shield status outright. The safe harbor lives in one place only: the marketing literature of the age assurance industry, which has an obvious commercial interest in telling platforms that a certificate purchases peace.[66] The inversion resolves into something simpler than an escape hatch. No jurisdiction built a general one; the single genuine method defense in American positive law is the narrow, good faith conditioned provision of the Texas App Store Accountability Act, deliberately omitted from its Utah twin;[67] and an industry gloss has been standing in for law on three continents while the gloss's beneficiaries cite it to each other.

VII. Where Responsibility Lands

Remove the deflections and the liability map draws itself, because the responsibility settles where the knowledge sat all along.

---

65. *Mind the Gap*, *supra* note 25.
66. Klippa, *supra* note 2.
67. *Supra* Part VI.A.

A. The Platform

The platform faces the claim in its cleanest form: common law negligence for deploying and continuing to operate an access-control system it knew to be defeated by the population the system existed to stop, with the standard of care informed but not capped by the statutory mandate,[68] and negligence per se where a state enactment supplies the standard and the plaintiff-child sits within the protected class.[69] The protective-statute doctrine of Part IV.C strips the comparative-fault and misrepresentation defenses as to the child plaintiff; the very-hazard rule of Part V strips the superseding-cause defense as to both the child's evasion and the predator's purchased entry; Part VI strips compliance. What remains is the question negligence law was built to ask: whether a reasonable operator, knowing what this operator knew, on the day it knew it, would have run this gate this way. The five dollar listings, the published error bands, and the regulator's robustness criterion are the jury exhibits, and they are all the operator's own or its vendor's own documents.

B. The Age Assurance Vendor

The vendor has occupied the most comfortable seat in the architecture, invisible to the harmed child (no privity, no contact) and indemnified against the platform (by contract). Its exposure deserves more attention than the commentary has given it, and two theories carry it. The first sounds in negligent misrepresentation and its certification analogues: a vendor that markets a method as satisfying "highly effective" criteria, or that supplies conformity documentation a platform relies on for regulatory compliance, while its own published error bands and its own robustness testing predict child-level defeat rates at the threshold, has supplied information for the guidance of others in a transaction whose foreseeable downstream reliers include the regulator and, through the statutory scheme, the protected children themselves. The second

68. *See supra* Part VI.A.

69. *See* RESTATEMENT (SECOND) OF TORTS §§ 286, 288C; Vincent v. Riggi & Sons, Inc., 30 N.Y.2d 406 (1972) (statutory violation establishes fault as a matter of law); Britton v. Wooten, 817 S.W.2d 443 (Ky. 1991) (regulatory violations as negligence per se).

treats the verification system as a product: a system whose design accepts static images, resold sessions, or threshold-adjacent faces, when countermeasures were feasible and the regulator's own guidance named them (liveness detection, re-verification, challenge-age hardening),[70] presents an ordinary design-defect question, with the manufacturer's duty running to foreseeable users and foreseeable misusers alike.[71] The indemnity clauses that route vendor exposure back to platforms do not defeat either theory; they merely determine, between the two commercial parties, who funds the judgment, which is a matter of supreme indifference to the child holding it.

C. The State, in the Register It Can Be Reached In

The legislature cannot be sued, and this Article has not pretended otherwise. But the state is reachable in the register of design critique, and the doctrinal analysis above sharpens that critique into something more pointed than the familiar privacy objections. A mandate that names methods rather than outcomes manufactures the compliance narrative this Article spent Part VI dismantling; every method the statute names becomes a ritual whose performance the operator will plead, and the legislature's own enhanced penalties for the bypass event prove it knew the ritual would fail. The regimes that wrote outcome duties with circumvention-resistance criteria, whatever their other defects, at least declined to hand the industry its script. The deeper design question, whether the gate model itself, walling children out of spaces left dangerous, should yield to a duty model that makes the spaces safe for the users who will foreseeably be inside them, is beyond this Article's doctrinal scope, but the liability analysis gestures at its answer. Every doctrine surveyed here, the protective statutes, the very-hazard rule, the floor-not-ceiling principle, the robustness criteria, converges on one design philosophy: the party who builds the environment answers for the foreseeable occupants of the environment, including the ones who were not supposed to be there, precisely because their arrival was the anticipated event. A

70. Ofcom, *supra* note 14, ¶¶ 4.38–4.42 (challenge age calibration), 4.55–4.56 (liveness detection), 4.58 (repeated age checks to prevent "child access via device or account sharing").

71. *Cf.* Lemmon v. Snap, Inc., 995 F.3d 1085 (9th Cir. 2021) (manufacturers must "foresee all reasonable uses and misuses and the consequent foreseeable dangers" of their products).

legal architecture already this unanimous about who answers for the child inside the wall is an architecture quietly telling legislatures where the duty belongs.

VIII. Objections

A. "But She Lied"

The objection with the most intuitive force is the simplest: the child affirmatively misstated her age, and law should not reward lying. Three answers, in ascending order of finality. First, the premise mischaracterizes most of the record; a made-up face is not a statement, an estimator's threshold error is the vendor's output, and a birthdate field designed to verify nothing is, as the COPPA scheme's whole existence attests, not reliance infrastructure.[72] Second, where the child did misstate, the misstatement lives inside an instrument that never formed and a reliance the statutes forbid, so it has no doctrinal vehicle.[73] Third, and conclusively, the protective-statute cases decided this exact objection against the regulated party a century ago, on facts stronger for the defendant than any platform will ever have: a sixteen year old who verbally claimed to be over eighteen, face to face, to a seller with no verification technology at all, and the seller was held responsible anyway, because the statute's purpose "would be defeated" by any other rule.[74] The platforms are asking for a defense the gun shops were denied, while holding verification tools the gun shops never had.

B. "Perfect Verification Is Impossible"

True, and beside the point twice over. The claim developed here has never been that a bypass occurred; it is that operators continued running gates whose specific, documented failure modes they declined to mitigate while the mitigations were feasible, named in guidance, and in some

72. *Cf.* 16 C.F.R. pt. 312 (operator obligations turn on what the operator knows and what the service is directed to, a scheme that would be pointless if a self-declared birthdate settled anything).

73. *See supra* Parts IV.A–B.

74. Tamiami Gun Shop v. Klein, 116 So. 2d 421 (Fla. 1959); Koester v. Rochester Candy Works, 194 N.Y. 92 (1909) ("No principle of estoppel is applicable."), *quoted in* Vincent v. Riggi & Sons, Inc., 30 N.Y.2d 406 (1972).

cases already deployed by competitors. Impossibility of perfection has never excused failure to take the reasonable measures short of it; that is the floor-not-ceiling rule restated.[75] And the regulators priced the objection into their standards: Ofcom demands mitigation of circumvention "easily accessible to children," not of all circumvention conceivable,[76] which is a reasonableness standard wearing regulatory clothes. A defendant who cannot stop a nation-state can be expected to notice its verification badges selling for five dollars on eBay.

C. "Compliance Must Mean Something"

It does. Compliance is competent evidence of due care and will defeat many claims on many facts; the ordinary case with no unusual circumstances may end there.[77] What compliance has never meant, in any of the three jurisdictions, is absolution in the face of specific knowledge of an operating failure mode. The unusual circumstance that removes the case from the ordinary run is knowledge, and the bypass record of Part II is a knowledge record. Operators who prefer the shield to be thicker have their remedy: make the gate work against the circumvention children actually use, which is all any of the three regimes ever asked.

D. "This Proves the Whole Project Is Futile"

A reader persuaded by everything above might conclude that age gates are theater and the mandates a mistake, and take this Article as the brief for abandonment. That is not its argument. The Article's argument is narrower and, for the child at the gate, more useful: whatever the wisdom of the mandates, the responsibility structure they generated has been misread by every actor with an interest in misreading it, and correctly read, the structure holds the builders to their walls. Whether legislatures should prefer duties that make the rooms safe over rituals at the doors is a design question this Article has flagged and left; what it has not left is the child, standing inside the wall she was never supposed to cross, being told the crossing was her fraud. She was

---

75. Miller v. Warren, 390 S.E.2d 207 (W. Va. 1990).

76. Ofcom, *supra* note 44.

77. Ramirez v. Plough, Inc., 6 Cal. 4th 539, 548 (1993); RESTATEMENT (SECOND) OF TORTS § 288C cmt.

the foreseeable child. Everyone was paid, mandated, or certified on the premise that she was coming. The law of three continents, read without the industry's gloss, agrees on what follows from that premise, and it is time the courts were asked.

IX. Conclusion

The question this Article opened with, who is responsible when a child defeats a mandated age gate and is harmed, turns out to have been answered piecemeal by doctrines that predate the internet by most of a century, and the deflection architecture survives only by keeping the doctrines apart. Contract law, asked alone, is diverted into voidability. Fraud, asked alone, presumes the reliance the statutes forbid. Superseding cause, asked alone, casts the child as the intervening stranger. Compliance, asked alone, is inflated from floor to ceiling by an industry with certificates to sell. Assembled, the doctrines close every exit. The child cannot be charged, by the express holding of the protective-statute cases and the founding premise of the verification regimes themselves. The bypass cannot supersede, because it is the hazard. The compliance shield, outside one narrow and good faith conditioned Texas provision, does not exist in the texts of any of the three jurisdictions that supposedly grant it. What remains standing at the gate, when the scripts are read against their sources, are the parties who built it, certified it, and kept it running past the day its failure went on sale for five dollars: the parties who knew. Responsibility, in the end, is not deflected by the architecture. It is only delayed by it, and the delay ends the first time a court is asked the formation question, the very-hazard question, and the robustness question in the same complaint.