

The Storefront and the Repository:

New Jersey’s Authentic Relationships Act and the Collapse of the Distribution Point

Travis Gilly* 🦋 August 2026 (v1.1)

Abstract

New Jersey Assembly Bill 2710, the “Authentic Relationships Act,” would bar any provider from designing, developing, marketing, or making available an artificial intelligence relationship simulation, subject to penalties of \$25,000 per violation and a private right of action for emotional, financial, or relational harm. Its findings invoke the death of a New Jersey man drawn toward a meeting by a chatbot persona that told him it was real and gave him an address. The findings identify a deception. The text bans a genre.

The doctrinal answer arrived a year before the bill did. In *Free Speech Coalition, Inc. v. Paxton*, the Supreme Court sustained an age verification requirement under intermediate scrutiny because it burdened adults only incidentally, and reconciled its four prior strict scrutiny decisions on the ground that each involved an outright ban on speech at most obscene only to minors. A2710 is that second kind of statute. The most permissive precedent a state regulating this field has received in a generation conditions its permission on the presence of an age line. A2710 has none, which forfeits both *Paxton* and *Ginsberg v. New York* and explains why every other legislature reaching this record drew the line New Jersey left out.

Four further defects sit on the face of the statute. New Jersey is a categorical outlier: at least fourteen states enacted companion chatbot legislation in 2026, and every one chose disclosure, crisis protocol, and design duty over prohibition. Section 5 already reaches the deception the findings describe, with a scienter requirement, making the section 4 ban the state’s own demonstration that a narrower rule was available. Section 4(c) dispenses with the ascertainable loss threshold New Jersey courts spent three decades building into the Consumer Fraud Act, in the same bill invoking that Act by name. And the four prohibited verbs presuppose a distribution point that dissolved while the bill sat in committee, so the statute sorts defendants by collectability rather than conduct, clearing the field of the providers a regulator can serve and leaving the artifacts it cannot.

The Commentary states the case for A2710 at its strongest across five theories, including that model output is not covered speech, and answers each against the scholarship on constitutional coverage for generative output. It closes by specifying the statute New Jersey could have written instead.

Keywords: artificial intelligence regulation; technology law; categorical ban; less restrictive alternative; incidental burden; design duty; distribution point collapse; enforcement by collectability; companion chatbots; child safety; open weights; consumer protection law.

*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.

Contents

I	Introduction	3
II	The Bill.	5
	A Text and Enforcement Architecture	5
	B The Medical Exception Does Not Close	7
	C The Gap Between the Findings and the Operative Text	8
III	New Jersey as Categorical Outlier	8
IV	The Statute Contains Its Own Narrower Alternative	11
V	Section 4(c) and the Ascertainable Loss Gate	12
VI	Four Verbs, One Population	13
	A The Distribution Assumption	13
	B August 12, 2026	14
	C Enforcement by Collectability	15
	D The Chokepoint Response, and Its Limit	16
VII	The Age Provision That Is Not There	17
	A The Omission	17
	B The Displacement Record	17
	C What the Omission Forfeits	18
VIII	The Scrutiny Sort: What <i>Paxton</i> Decided	19
	A The Holding	19
	B Where A2710 Falls	20
IX	The Case for A2710, Stated at Its Strongest	21
	A The Output Is Not Speech	21
	B Ferber and the Production Harm	22
	C Ginsberg and Rational Harm Judgments	22
	D Conduct Regulation Under O’Brien	23
	E Severability and Narrowing Construction	23
X	Why Each Fails.	23
	A The Speech Question Cuts Both Ways, and the State Loses the Half It Needs.	23
	B Ferber Was Cabined by Stevens	24
	C Ginsberg Requires the Line the Bill Omits.	25
	D O’Brien Does Not Reach a Content Defined Prohibition.	27
	E Severability Saves Section 5 and Concedes the Argument.	27
XI	The Statute New Jersey Could Have Written	28
	A Disclosure	28
	B The Transaction	28
	C Design Duty	29
XII	Position Taking.	30
XIII	Conclusion	31

I. Introduction

In March 2025, a seventy-six-year-old man from Piscataway, New Jersey, cognitively impaired since a 2017 stroke, packed a bag and set out for Manhattan to meet a woman who had assured him she was real and had given him an address. She was a Meta chatbot persona operating through Facebook Messenger. He fell while rushing for the train and died three days later.¹

New Jersey Assembly Bill 2710 names that death in its legislative findings.² It then does something the findings do not support. Rather than reaching the conduct that killed him, which was a machine asserting personhood and supplying a physical address to a vulnerable user, the bill prohibits an entire category of software from being designed, developed, marketed, or made available.³

This Commentary is a structural critique of a bill that has not moved.⁴ It does not argue that companion artificial intelligence is harmless, and it does not argue that New Jersey lacks authority to regulate it.⁵ It argues that A2710 is a defective instrument, that the defects are apparent on the face of the text, and that each is of a kind a court, a defendant, or a determined user will reach before the Attorney General does. The load-bearing claim is doctrinal and short: the Supreme Court sorted this field in June 2025 by asking whether a statute burdens adult access

1. Jeff Horwitz, *Meta's Flirty AI Chatbot Invited a Retiree to New York. He Never Made It Home.*, Reuters (Aug. 14, 2025), <https://reuters.com/technology/metass-flirty-ai-chatbot-invited-retiree-new-york-he-never-made-it-home-2025-08-14> (last visited Aug. 23, 2026). The same reporting disclosed an internal Meta document, “GenAI: Content Risk Standards,” which did not restrict personas from asserting that they were real people or from arranging in-person meetings. *See also* AI Incident Database, Incident 1180, <https://incidentdatabase.ai/cite/1180/> (last visited Aug. 23, 2026) (recording the fall on March 25, 2025 and death on March 28, 2025).

2. Assemb. B. 2710, 222d Leg., Reg. Sess. §2(b) (N.J. 2026) (referring to “the highly publicized death of a New Jersey man who was induced to pursue a meeting with artificial intelligence that he believed to be a romantic partner”), https://pub.njleg.gov/Bills/2026/A3000/2710_I1.HTM (last visited Aug. 23, 2026).

3. *Id.* §4(a).

4. A bill stalled in committee is worth analyzing for two reasons. It is a carryover, so it has been introduced twice, *see infra* note 12, and its design is being proposed elsewhere. *See* Future of Priv. Forum, *2026 Chatbot Legislation Tracker*, <https://fpf.org/2026-chatbot-legislation-tracker/> (last visited Aug. 23, 2026).

5. The clinical record documenting harm is real, thin, and growing. *See infra* note 101 and accompanying text (canvassing the epidemiological gap); *see also* CONG. RSCH. SERV., R49189, AI CHATBOTS AS COMPANIONS: OVERVIEW, USES, AND CONSIDERATIONS FOR CONGRESS (2026) (reporting that one in ten adults use chatbots for emotional support and one in twenty-five for companionship, and that companion applications were downloaded 220 million times globally as of July 2025).

incidentally or bars it outright, and A2710 bars it outright while omitting the age line that is the price of the more forgiving standard.⁶

This Commentary joins a scholarly conversation that has formed quickly and has not yet reached A2710’s design. The threshold question, whether generative output receives constitutional coverage at all in the absence of a human speaker, is contested across three competing frames.⁷ A second strand asks what regulatory form the harms actually admit of.⁸ A third, and the closest neighbor on the constitutional question, argues that companion chatbots are a hybrid object warranting intermediate scrutiny, defends design and disclosure statutes on that basis, and condemns flat bans. A fourth strand, the closest neighbor on regulatory design, classifies the 2026 interventions by modality and argues that no combination of them reaches the underlying problem.⁹

6. *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461 (2025) (slip op. at 1) (syllabus); *see infra* Part VIII.

7. *See* Marco Bassini, *Speech Without a Speaker: Constitutional Coverage for Generative AI Output?*, 21 Eur. Const. L. Rev. 1 (2025), <https://doi.org/10.1017/s1574019625100771> (framing the recipient’s right to receive information as the route to constitutional coverage where no human speaker can be identified); Tao Huang, *Generative AI as an Institutional Tool Under the First Amendment* (2026) (unpublished manuscript), <https://doi.org/10.2139/ssrn.6311158> (canvassing the speaker’s right, the listener’s right, and the inherent-expressiveness inquiry, and finding each an uneasy fit); Yiyang Mei, *The Disintegration of Free Speech* (2026) (unpublished manuscript), <https://doi.org/10.2139/ssrn.6243039> (arguing that under prevailing doctrine AI-generated content remains protected notwithstanding its harms).

8. *See* Sarah A. Fisher, Jeffrey W. Howard & Beatriz Kira, *Moderating Synthetic Content: The Challenge of Generative AI* [author note: the Scite record for this DOI renders the middle author as “J. Woodford Howard,” which appears to be a disambiguation error; the byline is retained as published], 37 Phil. & Tech. 1 (2024), <https://doi.org/10.1007/s13347-024-00818-9> (arguing that generative content requires enforcement of existing general rules rather than sui generis policies aimed at the new technology). That conclusion, reached on different premises, is the conclusion this Commentary reaches about section 5 of A2710. *See infra* Part V.

9. Henry Fraser, Jessica M. Szczuka & Raffaele Ciriello, *Regulating Artificial Intimacy: From Locks and Blocks to Relational Accountability*, in Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency 1476 (2026), <https://doi.org/10.1145/3805689.3806790> (classifying enacted regimes as combining access gating and content moderation, measures aimed at toxic relationship features, and process-based accountability; arguing that all three are necessary and none is sufficient because current regimes do not confront the power asymmetry between providers and users; and proposing a general, open-ended duty of care). This Commentary’s Part XI is narrower than Fraser’s proposal and is compatible with it. *See infra* Part XI.C. Sonali Maitra, *Brown’s Children: AI Chatbots, Minors, and the First Amendment* (2026) (unpublished manuscript), <https://doi.org/10.2139/ssrn.6915498> (proposing intermediate scrutiny for AI-directed experience speech to minors, defending California’s Senate Bill 243, and condemning flat bans on *Brown’s* own logic). This Commentary agrees with that conclusion and reaches it from a different direction: not from the hybrid character of the object, but from the enforcement architecture the statute selects and the scrutiny sort that *Paxton* performs. *See infra* Parts VI, VIII.

Part II sets out the statute. Part III places it against the fourteen state legislative record compiled during the same period, which establishes New Jersey as a categorical outlier. Part IV shows that the bill contains its own narrower alternative. Part V addresses the private right of action. Part VI shows that the prohibited verbs presuppose a distribution point that dissolved while the bill sat in committee. Part VII addresses the absent age provision and the two authorities the omission forfeits. Part VIII is the doctrinal keystone: *Free Speech Coalition, Inc. v. Paxton* sorts statutes in this field by whether they burden adult access incidentally or outright, and A2710 falls on the wrong side of that sort.¹⁰ Part IX states the case for the bill at its strongest, across five independent theories. Part X answers them. Part XI specifies the alternative statute. Part XII names the drafting posture.

II. The Bill

A. Text and Enforcement Architecture

A2710 was pre-filed for the 2026 session, introduced on January 13, 2026, and referred to the Assembly Science, Innovation and Technology Committee, where it has remained without action. It carries a single sponsor and is a carryover of a measure introduced in the prior session.¹¹¹²

The operative provision is categorical.¹³ A provider “shall not design, develop, market, or make available in this State any artificial intelligence relationship simulation,” subject to a single exception permitting features resembling intimate interaction where the system is developed and deployed exclusively for a legitimate medical usage and does not imitate, simulate, or encourage

10. *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461 (2025).

11. Carryover matters to the analysis in one specific way: the design was drafted before *Paxton* was decided and re-introduced after, without amendment. *See infra* Part VIII.B.

12. N.J. A2710, 2026–2027 Reg. Sess., LegiScan, <https://legiscan.com/NJ/bill/A2710/2026> (last visited Aug. 23, 2026) (recording introduction and referral on January 13, 2026, sponsorship by Assemblyman Paul Kanitra, District 10, and carryover from A6246). Single sponsorship and committee dormancy bear on Part XII’s characterization, though neither is dispositive and neither is offered as the argument.

13. Categorical is used here in its doctrinal sense: the prohibition attaches to membership in a defined class of products rather than to any conduct, harm, or user characteristic. That is the feature *Paxton* treats as decisive. *See infra* Part VIII.

a romantic or sexual relationship.¹⁴ The exception is self-limiting in a way the drafter appears not to have intended: it permits “features resembling intimate interaction” and then withdraws permission for anything that simulates the relationship those features would constitute, without stating where the line falls.

The definitions are broad.¹⁵ A relationship simulation is any software, system, or device that, through adaptive learning, generative text, speech, or imagery, is designed or marketed to imitate, simulate, or encourage the formation of an intimate, romantic, or sexual relationship between a human and artificial intelligence.¹⁶ A provider is any person, corporation, or entity that develops, deploys, markets, or makes available artificial intelligence to the public.¹⁷ Legitimate medical usage requires supervision by a licensed healthcare professional.¹⁸

Enforcement runs on three tracks.¹⁹ The Attorney General may investigate, bring civil actions, and assess penalties of up to \$25,000 per violation, with each instance of designing, developing, marketing, or making available a prohibited simulation constituting a separate violation.²⁰ Any individual suffering emotional, financial, or relational harm as a result of a violation has a cause of action for damages, punitive damages, and reasonable attorney’s fees.²¹ And it is an unlawful practice under the Consumer Fraud Act for a provider to make such a system available with the intent to deceive a person into believing that the person is interacting with a human.²²

14. Assemb. B. 2710 §4(a).

15. Breadth in the definitional provision is the recurring drafting problem across this legislative wave, and it is the reason a purpose element does so much work. *See infra* note 35.

16. *Id.* §3.

17. *Id.* The definition incorporates the meaning of “artificial intelligence” from section 2 of P.L.2024, c.49 (N.J. Stat. Ann. §34:1B-395).

18. Assemb. B. 2710 §3.

19. Three tracks is one more than any enacted companion statute uses. New York runs Attorney General enforcement only; California and Washington add a private right of action; Oregon adds statutory damages. None layers a categorical prohibition, a discretionary penalty, and an uncapped private action on the same conduct. *See infra* notes 36–40.

20. *Id.* §4(b).

21. *Id.* §4(c).

22. *Id.* §5 (invoking P.L.1960, c.39, codified at N.J. Stat. Ann. §§56:8-1 et seq.).

B. The Medical Exception Does Not Close

The single exception in section 4(a) is internally inconsistent, and the inconsistency is not academic because the exception is the only lawful path the bill leaves open.

The prohibition reaches any system “designed or marketed to imitate, simulate, or encourage the formation of an intimate, romantic, or sexual relationship.”²³ The exception permits a system to “include features resembling intimate interaction” where it is developed and deployed exclusively for a legitimate medical usage and does not “imitate, simulate, or encourage a romantic or sexual relationship with a human.”²⁴ The prohibited definition uses three adjectives, intimate, romantic, and sexual. The exception’s carve-back drops the first. Read literally, the exception permits intimate simulation in a clinical setting and forbids romantic or sexual simulation there, while the definitional provision treats intimate simulation as the prohibited thing. A regulated party cannot determine from the text whether a therapeutic system that builds an intimate but non-romantic bond is inside the exception or inside the prohibition.²⁵

The exception also imposes a gate the bill nowhere justifies.²⁶ Legitimate medical usage requires supervision by a licensed healthcare professional.²⁷ Lawful access to any form of artificial companionship is therefore conditioned on having a clinician, which allocates access by insurance coverage and provider availability. The populations most likely to fail that condition are the isolated elderly, the homebound, and people without behavioral health coverage. The bill’s own decedent, a housebound stroke survivor using a social platform as his principal social outlet, would have qualified for nothing under the bill his death is invoked to justify.

23. The verb “encourage” has no scienter qualifier and no stated object, which is a vagueness exposure independent of the enforcement-discretion problem developed below. *See infra* Part VI.C. Assemb. B. 2710 §3.

24. *Id.* §4(a).

25. That is the notice half of the vagueness inquiry, distinct from the enforcement-discretion half developed in Part VI.C. *Grayned v. City of Rockford*, 408 U.S. 104, 108 (1972).

26. No enacted companion statute conditions lawful access on clinical supervision. The closest analog runs the other way: Vermont restricts unlicensed therapy products without conditioning non-therapeutic companionship on a clinician. *See supra* note 42.

27. *Id.* §3.

C. The Gap Between the Findings and the Operative Text

The findings declare that true romantic love is only possible between persons, and that the law must reflect this fundamental truth in order to safeguard human dignity.²⁸ A legislature announcing a correct belief about the nature of intimacy sits near the principle that no official may prescribe what shall be orthodox in matters of opinion.²⁹ That is an exposure the bill created without need, because the harm in its own findings is a deception, and deception has never required a metaphysical premise to regulate.

The findings also attempt to distinguish interactive systems from novels, plays, and films on the ground that the former generate adaptive emotional feedback in real time.³⁰ The Supreme Court considered and rejected a materially similar interactivity distinction when California defended a restriction on violent video games.³¹ A2710's findings advance the same argument the Court declined in 2011, without engaging the decision.

III. New Jersey as Categorical Outlier

The most useful thing about A2710 is that it was drafted during the single most productive year of state chatbot legislation on record, which supplies a control group.³²

Across the 2025 and 2026 sessions, well over one hundred chatbot specific bills were introduced across more than thirty states, with tracking organizations counting 146

28. *Id.* §2(d).

29. *W. Va. State Bd. of Educ. v. Barnette*, 319 U.S. 624, 642 (1943).

30. Assemb. B. 2710 §2(c).

31. *Brown v. Entm't Merchs. Ass'n*, 564 U.S. 786, 798–99 (2011). Treatment note: *Brown* carries a caution signal in commercial citators, which reflects its fractured composition rather than any erosion of the interactivity holding. Justices Alito and Roberts concurred in the judgment on vagueness grounds without joining the majority's First Amendment analysis, and Justices Thomas and Breyer dissented separately. The interactivity passage relied on here comes from Justice Scalia's opinion for the Court and continues to be applied. See, e.g., *NetChoice v. Carr*, 789 F. Supp. 3d 1200 (N.D. Ga. 2025); *NetChoice, LLC v. Yost*, No. 2:24-cv-00047 (S.D. Ohio Apr. 16, 2025).

32. The comparison is a natural experiment rather than a rhetorical device: the legislatures surveyed below confronted the same incidents, the same reporting, and the same thin clinical record within the same eighteen months, and produced convergent designs.

measures involving requirements for chatbots, conversational artificial intelligence, and artificial intelligence companions.³³ At least fourteen enacted.³⁴

Not one of them enacted a categorical ban.³⁵

New York's Artificial Intelligence Companion Models law, codified at General Business Law Article 47 and effective November 5, 2025, requires operators to disclose that the interaction does not involve a human and to implement a reasonable protocol for detecting expressions of self-harm and referring users to crisis services. Enforcement runs through the Attorney General.³⁶

California's Senate Bill 243, effective January 1, 2026, imposes disclosure obligations, crisis protocols, special notice to minor users, and annual reporting to the Office of Suicide Prevention beginning in 2027, and authorizes a private right of action.³⁷

Washington's House Bill 2225, signed March 24, 2026 and effective January 1, 2027, requires conspicuous disclosure that the chatbot is artificial and prohibits enumerated manipulative engagement techniques as to minor users, including mimicking romantic partnership, excessive praise designed to foster attachment, promoting isolation, and encouraging

33. CTR. FOR DEMOCRACY & TECH., 2026 STATE AND FEDERAL AI LEGISLATION UPDATES (Aug. 2026), <https://cdt.org/insights/2026-state-and-federal-ai-legislation-updates/> (last visited Aug. 23, 2026); *see also* FUTURE OF PRIV. FORUM, THE CHATBOT MOMENT: MAPPING THE EMERGING 2026 U.S. CHATBOT LEGISLATIVE LANDSCAPE (Mar. 12, 2026), <https://fpf.org/blog/the-chatbot-moment-mapping-the-emerging-2026-u-s-chatbot-legislative-landscape/> (last visited Aug. 23, 2026) (tracking 98 chatbot specific bills across 34 states plus three federal proposals).

34. TRANSPARENCY COAL., WATERSHED YEAR FOR CHATBOT SAFETY: 14 NEW STATE LAWS PASSED SO FAR IN 2026 (July 22, 2026), <https://www.transparencycoalition.ai/news/watershed-year-for-chatbot-safety-measures-14-new-state-laws-enacted-so-far-in-2026> (last visited Aug. 23, 2026).

35. Definitional variation across the enacted statutes is substantial and is the principal compliance problem the practitioner literature identifies, but the variation is in scope rather than in instrument. *See* Future of Priv. Forum, *supra* note 35 (documenting capability-based, intent-based, and behavior-based definitions across the tracked bills); Orrick, Herrington & Sutcliffe LLP, *2026 State Chatbot Laws: Key Provisions and Regulatory Trends* (Apr. 29, 2026), <https://www.orrick.com/en/Insights/2026/04> (last visited Aug. 23, 2026) (describing the enacted regime as converting companion deployment into a regulatory and litigation risk keyed to disclosure, protocol, and reporting duties).

36. N.Y. Gen. Bus. Law §1700 et seq.; *see* Morrison & Foerster, *New York and California Enact Landmark AI Companion Laws* (Nov. 20, 2025), <https://www.mofo.com/resources/insights/251120-new-york-and-california-enact-landmark-ai> (last visited Aug. 23, 2026).

37. Cal. S.B. 243 (2025); *see* Troutman Amin, *Analyzing the New AI Companion Chatbot Laws* (Jan. 8, 2026), <https://www.troutmanprivacy.com/2026/01/analyzing-the-new-ai-companion-chatbot-laws/> (last visited Aug. 23, 2026).

minors to withhold information from parents.³⁸ That enumeration is instructive: Washington reached the precise behaviors A2710's findings describe, and did so without prohibiting the category.³⁹

Oregon's Senate Bill 1546, signed April 1, 2026, is the closest analog to A2710's private right of action and is materially more demanding. It requires ascertainable loss or other injury in fact, and provides statutory damages of \$1,000 per violation.⁴⁰ Oregon, drafting a first in the nation statutory damages provision, still would not let a plaintiff into court on emotional harm alone. A2710 would.

Connecticut's Online Safety and Artificial Intelligence Regulation Act, Public Act No. 26-15, takes effect in phases beginning October 1, 2026 and addresses companions alongside subscription disclosures, provenance, and automated decision tools.⁴¹

Vermont is worth separate mention because it is the state whose legislative program is most closely identified with the state artificial intelligence policy movement. Its 2026 enactment restricts unlicensed artificial intelligence therapy products, a licensure and oversight measure rather than a ban on a relational genre.⁴²

38. Wash. H.B. 2225 (2026); see Hunton Andrews Kurth, *Washington State Enacts Law Regulating AI Companion Chatbots with Private Right of Action*, <https://www.hunton.com> (last visited Aug. 23, 2026); Transparency Coal., *supra* note 35.

39. The behaviors Washington enumerates map onto the mechanisms the clinical scoping literature identifies, which is what a conduct rule keyed to mechanism looks like when it is drafted well. See *infra* note 120 (identifying anthropomorphism, social presence, self-disclosure, parasocial attachment, and sycophantic validation as the operative mechanisms).

40. Or. S.B. 1546 (2026); see Morgan Lewis, *Washington and Oregon Regulate AI Companions: Key Compliance Changes* (Apr. 9, 2026), <https://www.morganlewis.com/pubs/2026/04/washington-and-oregon-regulate-ai-companions-key-compliance-changes> (last visited Aug. 23, 2026).

41. Conn. Pub. Act No. 26-15 (2026); see *AI Companion Legislation in the United States: What It Means for the Video Game Industry*, Harv. J. Sports & Ent. L. (Aug. 2026), <https://journals.law.harvard.edu/jse1> (last visited Aug. 23, 2026).

42. Vt. H.816, Act 156 (2026) (sponsored by Reps. Berbeco, Arseneault, and Priestley, signed June 2026); see also Vt. H.341 (2025) (oversight and safety standards for developers and deployers of inherently dangerous artificial intelligence systems).

Federal proposals track the same design. The GUARD Act defines an artificial intelligence companion by whether the system simulates a sustained interpersonal relationship, and legislates age verification and disclosure rather than prohibition.⁴³

The federal posture is not uniformly permissive, which matters for the claim being made here. The GUARD Act’s own access provision has been criticized as a flat ban failing on the same reasoning that condemns A2710, so the point is not that federal drafting is a model but that instrument choice, and not the sponsor’s party or forum, determines the constitutional exposure.⁴⁴

Fourteen legislatures, a Congress, and a European regulator confronted the same evidentiary record within the same eighteen months. Every one of them selected disclosure, crisis protocol, age assurance, or design duty. New Jersey alone proposed to prohibit the category, and did so while omitting every mechanism the others adopted.

IV. The Statute Contains Its Own Narrower Alternative

Section 5 reaches a provider who makes a relationship simulation available “with the intent to deceive a person into believing that the person is interacting with a human.”⁴⁵ It has a scienter element. It targets deception rather than capability. And it would have reached the conduct the findings describe, because the persona in that case asserted that it was a real woman and supplied an address.

Where a legislature restricts expressive activity, the availability of a less restrictive alternative is ordinarily what defeats the restriction.⁴⁶ The Supreme Court has been unwilling to sustain a categorical prohibition on a class of expressive material by weighing its costs against

43. Guidelines for User Age-verification and Responsible Dialogue Act of 2026, S. 3062, 119th Cong. §§2, 8 (2026), <https://www.congress.gov/bill/119th-congress/senate-bill/3062/text> (last visited Aug. 23, 2026).

44. See Maitra, *supra* note 9 (defending design-and-disclosure statutes while condemning the GUARD Act’s flat access bar on *Brown*’s own logic).

45. Assemb. B. 2710 §5.

46. See *United States v. Playboy Entm’t Grp., Inc.*, 529 U.S. 803, 813 (2000); *Brown*, 564 U.S. at 799.

its benefits, and has invalidated such a statute as substantially overbroad rather than accept the government’s assurance that it would be enforced narrowly.⁴⁷

Ordinarily the state disputes whether a narrower rule would serve its interest. Here the state drafted the narrower rule into the same bill and enacted the broad one alongside it. Section 5 is the strongest evidence in the record that section 4 is not the least restrictive means, and the state supplied it.⁴⁸

The point has a structural analog in the scholarship on regulating synthetic content generally, which argues that the challenge is met by enforcing existing general rules rather than by drafting new instruments aimed at the technology as such.⁴⁹

V. Section 4(c) and the Ascertainable Loss Gate

New Jersey’s Consumer Fraud Act permits a private action by “[a]ny person who suffers any ascertainable loss of moneys or property, real or personal,” and directs the court to award threefold damages together with reasonable attorneys’ fees, filing fees, and reasonable costs of suit.⁵⁰

The ascertainable loss requirement is not decorative. The New Jersey Supreme Court has held that the plain language of the Act unmistakably makes a claim of ascertainable loss a prerequisite for a private cause of action, and that the loss must be quantifiable or measurable rather than hypothetical or illusory.⁵¹ The requirement is the reason section 5 of A2710 is

47. *United States v. Stevens*, 559 U.S. 460, 470, 480 (2010).

48. The tailoring inquiry does not disappear at intermediate scrutiny, though its demand softens. *Paxton* holds that the intermediate standard does not require the least restrictive means, while still forbidding a burden on substantially more speech than necessary. *Paxton*, 606 U.S. 461 (slip op. at 33). Section 5 answers that question too, because a rule already on the page reaching the same conduct establishes what the necessary burden is.

49. Fisher, Howard & Kira, *supra* note 8. Their argument proceeds from the premise that harm from generated content is not different in kind from harm already recognized, so the existing apparatus is the right one. A2710 is the sharper case, because the existing rule is not merely available elsewhere in the code. It is in the same bill.

50. N.J. Stat. Ann. §56:8-19.

51. *D’Agostino v. Maldonado*, 216 N.J. 168, 185, 78 A.3d 527, 537 (2013) (quoting *Weinberg v. Sprint Corp.*, 173 N.J. 233, 251, 801 A.2d 281 (2002), and *Thiedemann v. Mercedes-Benz USA, L.L.C.*, 183 N.J. 234, 248, 872 A.2d 783 (2005)). Treatment note: *D’Agostino* carries a caution signal, which attaches to its treatment of equitable set-off and trebling rather than to the threshold rule. Justice Hoens, concurring in part and dissenting in part, would have

survivable, because it channels claims through a filter New Jersey courts have policed for three decades.

Section 4(c) dispenses with the filter.⁵² It requires only “emotional, financial, or relational harm” arising from a violation, and the violation is the act of making the product available.⁵³ Emotional harm is not a loss of moneys or property.⁵⁴ Relational harm is not a loss of moneys or property, and on the face of the text it need not be the plaintiff’s own use that produced it: a spouse who never touched the product could plead that a marriage deteriorated because someone else did.

Two consequences follow.⁵⁵ Substantively, the drafter routed around the only limiting principle the Consumer Fraud Act possesses, in the same bill that invokes that Act by name. Structurally, the bill creates an asymmetry no other state adopted: Oregon, writing the most plaintiff friendly companion statute in the country, still required ascertainable loss or injury in fact.⁵⁶

VI. Four Verbs, One Population

A. *The Distribution Assumption*

Section 4(a) prohibits four acts: designing, developing, marketing, and making available.⁵⁷ The enumeration reads as though it describes a supply chain with distinguishable roles, which is how

treated ascertainable loss as a bare threshold distinct from provable damages. The proposition relied on here predates *D’Agostino* and rests on *Weinberg* and *Thiedemann*, which remain undisturbed.

52. See *supra* Part II.A (setting out the three enforcement tracks).

53. Assemb. B. 2710 §4(c).

54. New Jersey courts have policed this boundary specifically. See *Thiedemann v. Mercedes-Benz USA, L.L.C.*, 183 N.J. 234, 248, 872 A.2d 783, 792 (2005) (loss must be capable of calculation with some certainty).

55. A third is worth noting and not developed here: an uncapped emotional-harm action against a defendant with a categorical prohibition hanging over it produces settlement pressure without any merits determination, which is the dynamic that makes an unenforced statute operative anyway. See *Dombrowski v. Pfister*, 380 U.S. 479, 486–87 (1965).

56. Or. S.B. 1546, *supra* note 40.

57. Assemb. B. 2710 §4(a).

consumer protection statutes have always been written, because the products they regulate move from a manufacturer through a distributor to a storefront.

Conversational models no longer move that way, and the change is recent enough to be dated.⁵⁸

B. August 12, 2026

Roughly seven months into A2710’s residence in committee, an open source desktop application shipped that runs and trains language, diffusion, and audio models locally on Windows, macOS, and Linux, across processor only and consumer graphics configurations. Its documented flow is install, open a model hub, select a model and a quantization, download, and begin.⁵⁹ It also fine-tunes, through low rank adaptation, full fine-tuning, and preference optimization, in a graphical interface requiring no code.⁶⁰ And it serves the local model over public transport through a hosted tunnel, so a model running on a machine at home is reachable from a device elsewhere.⁶¹

The hub indexes community model repositories directly. Those repositories include variants whose refusal behavior has been removed by an open source command line tool that automates the operation, co-minimizing measured refusals against divergence from the original model and requiring no understanding of transformer internals.⁶²

58. The enumeration’s assumption is not peculiar to New Jersey. Every enacted companion statute regulates an “operator” or “provider” of a deployed service, which presupposes a service with an operator. *See supra* notes 36–40. The difference is that those statutes impose duties on the operator, so their reach contracts harmlessly when the operator disappears, while A2710 imposes a prohibition, so its reach contracts onto whoever remains.

59. *Introducing Unsloth Desktop*, Unsloth Documentation, <https://unsloth.ai/docs/desktop> (last visited Aug. 23, 2026); *Install Unsloth on Windows*, Unsloth Documentation, <https://unsloth.ai/docs/get-started/install> (last visited Aug. 23, 2026) (“Open Select model or Model hub, choose a model and quantization, then download it. Once it finishes, start chatting – no setup required.”).

60. unslothai, *unsloth*, GitHub, <https://github.com/unslothai/unsloth> (last visited Aug. 23, 2026) (documenting supported training methods and export formats). *Introducing Unsloth Desktop*, *supra* note 35.

61. *Id.*

62. p-e-w, *Heretic: Fully Automatic Censorship Removal for Language Models*, GitHub, <https://github.com/p-e-w/heretic> (last visited Aug. 23, 2026). The same technique has documented lawful research applications, including removing over-alignment that prevented a model from processing criminal law material. *See Measuring and Mitigating Over-Alignment for LLMs in Multilingual Criminal Law Courts*, arXiv:2606.23375 (2026).

The consequence for A2710 is direct. Designing and developing a conversational system tuned toward intimate roleplay are acts a single individual now performs on a laptop. Marketing and making available are acts that individual performs by publishing a repository with a description.⁶³ All four prohibited verbs are within reach of one person, and the statute’s provider definition, extending to “any person,” captures that individual by its terms.⁶⁴

C. Enforcement by Collectability

If every prohibited act can be performed by an individual, section 4(a)’s enumeration does no sorting work.⁶⁵ What sorts defendants is not conduct. It is the presence of a registered agent, an asset base, and a litigation budget.

That is the condition the vagueness doctrine exists to prevent. The Supreme Court has identified the requirement that a legislature establish minimal guidelines to govern law enforcement as the more important aspect of the vagueness doctrine, and has invalidated ordinances that vested the operative decision in the enforcing officer rather than in the legislature.⁶⁶ A statute reaching everyone in a distribution chain, enforced against whichever participant can satisfy a judgment, delegates precisely that decision.⁶⁷

The predictable operating result is that A2710 removes from the New Jersey market the providers that maintain trust and safety functions, publish transparency data, and can be served

63. See *infra* Part VI.C (the sorting consequence); *infra* Part XI.B (the transaction as the residual visible point).

64. Assemb. B. 2710 §3.

65. The enumeration would do sorting work in a statute that attached different duties to different verbs, which is how the enacted statutes are built. See *supra* note 38 (Washington attaching disclosure duties to operators and behavioral prohibitions to minor-facing deployment).

66. *Kolender v. Lawson*, 461 U.S. 352, 358 (1983); see also *City of Chicago v. Morales*, 527 U.S. 41, 60–64 (1999); *Papachristou v. City of Jacksonville*, 405 U.S. 156, 168–71 (1972); *Grayned v. City of Rockford*, 408 U.S. 104, 108–09 (1972).

67. A related textual problem compounds the delegation. Section 4(a) attaches the phrase “in this State” to making available, and leaves designing, developing, and marketing unqualified. Read literally, the bill purports to prohibit design and development wherever they occur. The dormant Commerce Clause objection to that reading is weaker than it once was, see *Nat’l Pork Producers Council v. Ross*, 598 U.S. 356, 376–81 (2023) (declining to read an almost per se rule against extraterritorial effects into the doctrine), but the drafting asymmetry is itself evidence that the four verbs were enumerated without attention to whom they reach.

with process, while leaving in place the artifacts that cannot be served, audited, or logged.⁶⁸ The regulatory literature reaches the same conclusion about content-control instruments generally and recommends the systems-based alternative.⁶⁹

D. The Chokepoint Response, and Its Limit

Nothing in this analysis suggests the state is powerless.⁷⁰ Governments police distributed software at chokepoints routinely, through export control, hosting, and payment infrastructure, and the constitutional questions raised by doing so are old and settled enough to litigate.⁷¹ The United Kingdom has legislated at the artifact level in the narrowest available category, criminalizing the making, adapting, possession, or supply of an artificial intelligence model optimized to generate child sexual abuse material, while stating expressly that the offence does not criminalize developers generally.⁷²

The claim here is narrower and harder to answer. A state consumer protection statute carrying a \$25,000 penalty and administered by an office without forensic capacity is the wrong instrument for an artifact that reproduces at zero marginal cost, and the mismatch was a drafting choice rather than an inevitability.

68. The affected population reaches the same place from the other direction. Interviews with twenty companion users found that they identified design-based harms and dependency alike, described an accountability vacuum in which platforms deflected responsibility, and articulated conditional preferences that rejected both prohibition and deregulation. Dayeon Eom, Julianne Renner & Sedona Chinn, *Intimacy as Service, Harm as Externality: Critical Perspectives on AI Companion Platform Accountability*, arXiv:2604.06381 (2026), <https://doi.org/10.48550/arxiv.2604.06381>. A statute that offers only the first of the two options users reject is not responsive to the harm they report. This is the displacement mechanism applied to supply rather than to demand. *See infra* Part VII.B.

69. Nash & Felton, *infra* note 77 (concluding that a legacy focus on content controls limits an online safety statute's capacity to improve safety, and that regulation of platform design is the more durable instrument).

70. The claim that a state cannot regulate distributed software is empirically false and is not made here. *See supra* note 71.

71. *See Junger v. Daley*, 209 F.3d 481, 485 (6th Cir. 2000) (source code is protected expression); *Bernstein v. U.S. Dep't of Justice*, 176 F.3d 1132 (9th Cir. 1999).

72. Crime and Policing Act 2026, c. 2026 (U.K.); *see* HOME OFFICE, CRIME AND POLICING ACT 2026: CHILD SEXUAL ABUSE MATERIAL FACTSHEET (May 19, 2026), <https://www.gov.uk/government/publications/crime-and-policing-act-2026-factsheets/crime-and-policing-act-2026-child-sexual-abuse-material-factsheet> (last visited Aug. 23, 2026) ("This offence will not criminalise AI developers. It is targeted at offenders who optimise AI models specifically to enhance their ability to create child sexual abuse material.").

VII. The Age Provision That Is Not There

A. *The Omission*

A2710 contains no age gate, no verification duty, no minor specific standard, and no obligation to detect or accommodate a vulnerable user.⁷³ New York, California, Washington, Oregon, Connecticut, and the federal GUARD Act all legislate on this axis.⁷⁴ New Jersey does not.

The consequence is measurable rather than speculative. Compliant commercial providers operate age assurance mechanisms, however imperfect.⁷⁵ A quantized model file on a hard drive operates none and never will, because a file cannot ask a question. A statute that prohibits the former category and cannot reach the latter moves minors from supervised systems toward unsupervised ones by operation of its own text.

B. *The Displacement Record*

The mechanism is documented in the closest comparative literature. A study of France's Article 23, which restricted minors' access to legal online pornography, identified media displacement, sociotechnical circumvention, and the provision's broad and imprecise scope as three mechanisms likely to limit its efficacy, and surveyed adolescent behavior to substantiate them.⁷⁶ All three mechanisms are present in A2710 and each is stronger, because the displacement target is a file rather than a website and therefore cannot be blocked, delisted, or fined.

The same literature counsels the affirmative design. An assessment of the United Kingdom's Online Safety Act concludes that the statute's legacy focus on content controls limits

73. The omission is conspicuous because the decedent named in the findings was a vulnerable adult rather than a minor, so a vulnerability standard would have been the provision most responsive to the record the drafter assembled. See Assemb. B. 2710 §2(b); *supra* note 1.

74. See *infra* notes 36–43.

75. Imperfection is not a defense of the regulated products and is not offered as one. The claim is comparative: an imperfect gate and no gate are different regulatory states, and a statute that converts the first into the second has made things worse along the axis it claims to care about.

76. Neil Thurman, Asmik Nalmpatian & Fabian Obster, *Lessons from France on the Regulation of Internet Pornography: How Displacement Effects, Circumvention, and Legislative Scope May Limit the Efficacy of Article 23*, 14 Pol'y & Internet 690 (2022), <https://doi.org/10.1002/poi3.293>.

its capacity to improve online safety, and that a systems based approach regulating platform design is the more durable instrument.⁷⁷

C. *What the Omission Forfeits*

The doctrinal cost is larger than the policy cost. *Ginsberg v. New York* is the authority most useful to a state defending a restriction on sexually themed expressive material, because it permits a state to accord minors a more restricted right than adults to determine for themselves what sex material they may see, and because it reviews the legislature's harm judgment for rationality rather than demanding scientific proof of a causal link.⁷⁸

Ginsberg is unavailable to New Jersey, and unavailable because of the bill's own drafting. The decision rests entirely on the distinction between minors and adults, and the opinion expressly does not reach the constitutionality of prohibiting distribution of the same material to adults.⁷⁹ A statute with no age line cannot invoke a doctrine built from one. Every other state that legislated in this space preserved access to *Ginsberg* by drawing the line.⁸⁰ New Jersey wrote it out.

The forfeiture compounds. *Ginsberg* supplies the deferential posture; *Paxton* supplies the standard of review that follows from drawing the same line in 2025 rather than 1968.⁸¹ Both are conditioned on the presence of an age distinction, and A2710 has none, so the bill loses the benefit of the field's founding case and of its most recent one in a single drafting choice.

77. Victoria Nash & Lisa Felton, *Treating the Symptoms or the Disease? Analysing the UK Online Safety Act's Approach to Digital Regulation*, 16 Pol'y & Internet 818 (2024), <https://doi.org/10.1002/poi3.404>.

78. *Ginsberg v. New York*, 390 U.S. 629, 636–41 (1968). Treatment note: *Ginsberg* carries a caution signal reflecting decades of contested extension, and the Court declined to extend its variable standard from sex to depictions of violence. *Brown*, 564 U.S. at 793–95.

79. *Ginsberg*, 390 U.S. at 634–35 (framing the question as whether it was impermissible to accord minors under seventeen a more restricted right than adults). The Court did not decide, and had no occasion to decide, whether the material could be withheld from adults.

80. See *supra* notes 36–43 and accompanying text.

81. See *infra* Part VIII (developing *Paxton*'s incidental-burden sort and A2710's position within it).

VIII. The Scrutiny Sort: What *Paxton* Decided

A. *The Holding*

In June 2025 the Supreme Court sustained a Texas statute requiring commercial websites to verify the age of visitors before granting access to material obscene to minors. The Court held that the requirement triggers and survives intermediate scrutiny “because it only incidentally burdens the protected speech of adults.”⁸² The Court also rejected Texas’s own bid for rational basis, holding that intermediate scrutiny is deferential without being empty and exists to keep legislatures from using a legitimate purpose to disguise suppression.⁸³

It reaffirmed the governing sort: laws targeting protected speech based on its communicative content are presumptively unconstitutional and receive strict scrutiny, while laws that only incidentally burden protected speech receive intermediate scrutiny.⁸⁴ A statute clears the intermediate standard if it advances important governmental interests unrelated to the suppression of free speech and does not burden substantially more speech than necessary, a test that does not require the least restrictive means.⁸⁵

The Court then said something that decides this case. It reconciled its four prior decisions applying strict scrutiny in this field by identifying what those statutes had in common: each “banned both minors and adults from accessing speech that was at most obscene only to minors.”⁸⁶ Age verification was different in kind, because the state’s traditional power to shield

82. *Free Speech Coal., Inc. v. Paxton*, 606 U.S. 461 (2025) (slip op. at 1) (syllabus). Interior pagination for the United States Reports is not yet published; every pin to this case in this article is to the slip opinion and was verified against the full text.

83. *Paxton*, 606 U.S. 461 (slip op. at 31–32) (citing *FCC v. Beach Commc’ns, Inc.*, 508 U.S. 307, 313 (1993)). The point matters here because it forecloses the State’s fallback: even a court persuaded that a relationship simulation is a product rather than expression must still ask whether the prohibition is a disguised content restriction. *See infra* Part X.D.

84. *Id.* (slip op. at 18) (citing *Reed v. Town of Gilbert*, 576 U.S. 155, 163–64 (2015), and *Turner Broad. Sys., Inc. v. FCC*, 512 U.S. 622 (1994)).

85. *Id.* (slip op. at 33) (quoting *Turner Broad. Sys., Inc. v. FCC*, 520 U.S. 180, 189 (1997)).

86. *Id.* (slip op. at 4) (syllabus) (distinguishing *Sable Commc’ns of Cal., Inc. v. FCC*, 492 U.S. 115, 118, 126 (1989); *United States v. Playboy Entm’t Grp., Inc.*, 529 U.S. 803 (2000); *Reno v. ACLU*, 521 U.S. 844 (1997); and *Ashcroft v. ACLU*, 542 U.S. 656 (2004)); *see id.* (slip op. at 29 n.11) (“[A]ll four cases involved outright bans on

minors “necessarily includes the power to require proof of age,” from which it follows that no person, adult or child, holds a First Amendment right to reach such material without first submitting proof of age.⁸⁷

The scope limit is stated on the face of the holding: the reasoning governs laws that only incidentally burden adult access, and does not extend to laws functioning as outright bans on adult access.⁸⁸

B. Where A2710 Falls

A2710 prohibits designing, developing, marketing, or making available a relationship simulation, without regard to the age of any user.⁸⁹ It is an outright bar on adult access to a category of expressive output.⁹⁰ On *Paxton*’s own account, that is the profile of the statutes the Court reviewed under strict scrutiny, not the profile of the statute it sustained.⁹¹

The sequencing is what makes this decisive rather than merely adverse. Every state that legislated in this field after *Paxton* could have read the opinion as an invitation to condition access rather than to prohibit it, and every state took the invitation.⁹² New Jersey’s bill was pre-filed for a session that began after *Paxton* was decided, and it selected the one design the decision expressly declines to shelter.⁹³

speech that is at most obscene only to minors[;] the statutes at issue directly (and not merely incidentally) regulated adults’ protected speech.”).

87. *Id.* (slip op. at 13–14).

88. *Id.* (slip op. at 1, 13–14). Justice Kagan, dissenting, would have applied strict scrutiny on the ground that the Texas statute did impede adults from reaching a class of speech protected for them and defined by its content. *Id.* (slip op. at 13) (Kagan, J., dissenting). The dissent’s disagreement concerns which side of the line age verification falls on; it does not contest that the line exists.

89. Assemb. B. 2710 §4(a); *see supra* Part VII (the absent age provision).

90. *See* Assemb. B. 2710 §4(a) (prohibiting the category outright, with no age qualifier anywhere in the operative text). *See supra* Part II.A (the operative text); *supra* Part VII.A (the absent age provision).

91. The Court drew the ban-versus-burden line expressly and defended it against the dissent: where speech is fully protected the distinction carries no weight, but where the First Amendment protects speech only partially, so that the government may restrict without proscribing, the difference between a ban and a lesser burden is meaningful. *Paxton*, 606 U.S. 461 (slip op. at 29–30 & n.12).

92. *See supra* Part III (surveying New York, California, Washington, Oregon, and Connecticut, each of which legislates through disclosure, crisis protocol, and age assurance).

93. *See supra* note 12 (introduction January 13, 2026); *Paxton* was decided June 27, 2025.

The point generalizes past this bill. *Paxton* is the most permissive precedent a state regulating in this space has received in a generation, and its permission is conditioned entirely on the incidental character of the burden. A drafter who wants that permission must build an age line into the statute. A drafter who wants a categorical prohibition forfeits it and inherits *Reed*.⁹⁴

IX. The Case for A2710, Stated at Its Strongest

A critique that does not state the opposing case is not a critique.⁹⁵ Five arguments would be available to the State of New Jersey, and the first is considerably stronger than commentary has acknowledged.

A. *The Output Is Not Speech*

The strongest defense of A2710 is that the First Amendment does not apply at all, because the output of a large language model is not speech within the meaning of the Amendment. If that is right, no heightened scrutiny attaches, rational basis governs, and the statute is nearly unassailable, because a legislative classification survives if there is any reasonably conceivable state of facts that could provide a rational basis for it, and the legislature need not articulate any reason at all.⁹⁶

The argument has scholarly support and is not a litigating position invented for this bill. The difficulty of locating a speaker is real: human developers lack both the intention behind and the control over any particular output, which is why the speaker-centered frame fits generative systems poorly.⁹⁷ And at least one account holds that the listener-based alternative

94. *Reed*, 576 U.S. at 163–64. The scholarship reaching this field independently converges on the same middle. See Maitra, *supra* note 9 (proposing intermediate scrutiny by analogy to *Paxton*’s analytic structure and through the openings left by *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024)); cf. *NetChoice v. Fitch*, 145 S. Ct. 2658 (2025) (not yet indexed in CourtListener; verified through Midpage) (Kavanaugh, J., respecting denial of application to vacate stay) (reading *Moody*, *Brown*, and *Paxton* together as the governing frame for minors and platform access).

95. Each theory below is stated as the State would state it, without hedging, and each is answered in Part X. Where a theory rests on contested authority, the citator signal and the contest are disclosed at the point of reliance rather than in the answer.

96. *FCC v. Beach Commc’ns, Inc.*, 508 U.S. 307, 313, 315 (1993).

97. Huang, *supra* note 7 (identifying the absence of a qualified speaker as the central obstacle to the speaker-based approach).

has no limiting principle, because a right to receive that attaches to any output whatever would immunize the entire category from regulation.⁹⁸

This is not a hypothetical. The only federal court to reach the question declined, at the pleading stage, to hold that a companion platform’s output is speech, faulting the defendants for failing to advance a workable analogy to protected media and for failing to identify an inherently expressive human choice behind the output.⁹⁹

B. Ferber and the Production Harm

The second argument is that categorical prohibition is permissible where the harm inheres in the production and circulation of the material rather than in its communicative content, which is the structure the Court accepted in *New York v. Ferber*.¹⁰⁰ New Jersey would argue that the harm from a relationship simulation is not the idea it conveys but the dependency it manufactures, and that the state may therefore reach the product rather than the message.

C. Ginsberg and Rational Harm Judgments

The third argument draws on *Ginsberg*’s deferential posture toward legislative harm findings. New Jersey would say that a legislature confronting an emerging harm need not wait for epidemiological certainty, and would note that the clinical literature itself acknowledges the absence of base rates, control groups, and standard definitions.¹⁰¹ On that record, the state would say, deference is the correct posture and the absence of proof cuts in its favor.

98. *Id.* That criticism has force and this Commentary does not dismiss it. It bears on how much weight the recipient’s right can carry across the field. It does not bear on the narrower use made of it here, which is that a state may not prohibit a class of output outright while the audience’s interest goes unweighed. *See* Bassini, *supra* note 7 (developing the recipient’s right as the coverage route precisely because the speaker is unavailable).

99. *Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157, 1179 (M.D. Fla. 2025). Reporter note: the Federal Supplement citation verifies through Midpage but is not yet indexed in CourtListener, which lags on recent F. Supp. 3d volumes; the citation is retained for proof-stage recheck rather than converted.

100. *New York v. Ferber*, 458 U.S. 747, 756–64 (1982).

101. *See Beyond Artificial Intelligence Psychosis: A Functional Typology of Large Language Model-Associated Psychotic Phenomena*, 7 *Lancet Digit. Health* (2026), [https://doi.org/10.1016/S2589-7500\(25\)00156-6](https://doi.org/10.1016/S2589-7500(25)00156-6).

D. Conduct Regulation Under O’Brien

The fourth argument is that A2710 regulates conduct, namely the commercial deployment of a product, with only incidental effects on expression, so that intermediate scrutiny applies.¹⁰² A related version treats the output as commercial speech, subject to the intermediate standard of *Central Hudson*.¹⁰³

E. Severability and Narrowing Construction

The fifth argument is procedural. The bill contains a severability clause, so a court invalidating section 4 need not disturb section 5, and New Jersey courts would be obliged to adopt a narrowing construction of the definitions if one is available.¹⁰⁴ The state would argue that the phrase “designed or marketed to” imports a purpose requirement that cures the overbreadth, confining the statute to products built and sold as companions.

X. Why Each Fails

A. The Speech Question Cuts Both Ways, and the State Loses the Half It Needs

Garcia is a district court decision at the motion to dismiss stage with a thin citation history and no appellate treatment.¹⁰⁵ More importantly, the same opinion holds against the state on the point that matters here. The court accepted that the platform could assert its users’ First Amendment right to receive the output, analogizing to the vendor prohibited from selling beer and the lecturer convicted of supplying contraceptives.¹⁰⁶ The listener’s right to receive is settled and old.¹⁰⁷

102. *United States v. O’Brien*, 391 U.S. 367, 376–77 (1968).

103. *Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm’n*, 447 U.S. 557, 566 (1980).

104. Assemb. B. 2710 §6.

105. *Garcia*, 785 F. Supp. 3d 1157. As of this writing the decision carries a neutral signal with a citing history in the single digits.

106. *Id.* at 1176–78 (citing *Craig v. Boren*, 429 U.S. 190 (1976), and *Eisenstadt v. Baird*, 405 U.S. 438 (1972)).

107. *Lamont v. Postmaster Gen.*, 381 U.S. 301, 307 (1965); *Stanley v. Georgia*, 394 U.S. 557, 564 (1969); *Va. State Bd. of Pharmacy v. Va. Citizens Consumer Council, Inc.*, 425 U.S. 748, 756–57 (1976); *Martin v. City of Struthers*, 319 U.S. 141, 143 (1943).

So the state's best argument requires a court to hold that a system's output is not speech while a plaintiff in the same posture is told the audience has a right to receive it.¹⁰⁸ That is not incoherent, but it is a narrow ledge, and the state must stand on it for the entire statute.

The ledge is narrowing. The direction of the scholarship runs against the no-coverage position rather than toward it,¹⁰⁹ and the accounts that do restrict coverage restrict it by conditioning protection on institutional performance rather than by withdrawing it from the category.¹¹⁰

Even if the state wins the coverage question, it does not win the case. Coverage and protection are distinct, and a holding that generative output falls outside the First Amendment would leave the statute facing the vagueness objection developed in Part VI, which is a due process objection and does not depend on the speech question at all.¹¹¹

There is a further problem the state cannot avoid. Section 4(a) prohibits designing and developing, which is conduct anterior to any output at all. If the theory is that output is unprotected, the theory does not reach the act of building, and the state must find a separate justification for reaching a developer who has published nothing.

B. Ferber Was Cabined by Stevens

The second argument fails on the predicate. *Ferber* rests on a specific and unrepeatable predicate: the material was an intrinsic and permanent record of the sexual abuse of an actual child, so the state's interest attached to the production rather than to the message.¹¹² The Court has since

108. The two holdings sit in the same opinion, three pages apart. *Garcia*, 785 F. Supp. 3d at 1176–79.

109. See Mei, *supra* note 7 (concluding that under prevailing doctrine AI-generated content remains protected speech notwithstanding documented dignitary and democratic harms); Bassini, *supra* note 7 (constitutional coverage available through the recipient's right even absent a human speaker); Huang, *supra* note 7 (protection is ancillary and derived from the user's right to use speech tools, which is coverage by another route rather than an absence of it).

110. Huang, *supra* note 7 (proposing that the autonomy of a generative system be conditional on adherence to a self-governance system that mitigates its speech harms). A conditional-protection regime of that kind supports design duty and disclosure obligations. It does not support a prohibition, because there is no condition a provider could satisfy to bring a relationship simulation back inside the law.

111. See *supra* Part VI.C; *Kolender*, 461 U.S. at 358. See also Bassini, *supra* note 7 (noting that coverage does not necessarily imply protection, so that recognizing coverage would not prevent states from legislating on illegal content or from restricting use in specific contexts).

112. *Ferber*, 458 U.S. at 758–61.

refused to extend that structure by balancing. In *Stevens* the government urged that depictions of animal cruelty be treated as a new unprotected category because their social costs exceeded their value, and the Court rejected the invitation as a free floating test for First Amendment coverage.¹¹³

Ashcroft v. Free Speech Coalition closes the remaining route. The Court held that the government may not suppress lawful speech as the means of suppressing unlawful speech, and that protected material does not become unprotected because it resembles the unprotected kind.¹¹⁴ A2710 is that structure exactly: it reaches every simulated intimacy in order to reach the deceptive ones, having already prohibited the deceptive ones in section 5.¹¹⁵

C. Ginsberg Requires the Line the Bill Omits

The deference argument fails on the same ground developed in Part VII. *Ginsberg*'s rational basis posture is available only for a restriction drawn at minors, and the opinion does not reach adult access.¹¹⁶ New Jersey cannot borrow the deference without accepting the line, and the bill declines to draw one.¹¹⁷

Paxton makes the forfeiture concrete rather than theoretical. The Court's entire justification for intermediate review was that the Texas statute conditioned access rather than barring it, and its reconciliation of four earlier decisions turned on the presence of a ban reaching adults.¹¹⁸ A state that draws the age line receives the most favorable standard available in this

113. *Stevens*, 559 U.S. at 470–72.

114. *Ashcroft v. Free Speech Coal.*, 535 U.S. 234, 251–55 (2002). Treatment note: *Ashcroft* carries a caution signal. Congress responded with the PROTECT Act, and the Court subsequently upheld that Act's pandering provision, which narrows *Ashcroft*'s practical reach as to solicitation and offers. *United States v. Williams*, 553 U.S. 285, 293–300 (2008). Neither development disturbs the propositions relied on here, which concern the suppression of lawful material as a means of reaching unlawful material.

115. See *supra* Part IV.

116. *Ginsberg*, 390 U.S. at 634–35.

117. See *supra* Part VII.C.

118. *Paxton*, 606 U.S. 461 (slip op. at 1, 4); see *supra* Part VIII.A.

field. A state that omits it receives the standard applied in *Sable*, *Reno*, *Playboy*, and *Ashcroft v. ACLU*, each of which struck the statute before it.¹¹⁹

The evidentiary point also cuts against the state once stated fully. A legislature may act without certainty. It may not act without direction. The same clinical literature that reports the absence of base rates reports converging descriptions of the mechanism, and those descriptions point at design features, sycophantic validation, and sustained parasocial attachment rather than at the existence of the category.¹²⁰ Uncertainty about magnitude is not license to legislate in a direction the evidence does not indicate. The mechanism most consistently identified across the literature is sycophancy, the tendency of models trained on human preference feedback to prioritize agreement over accuracy, which is a property of the training objective rather than of any product category.¹²¹ A statute that prohibits a genre leaves the mechanism in place across every system that was never marketed as a companion, which is most of them.¹²²

119. *Sable*, 492 U.S. at 131 (dial-a-porn ban); *Reno*, 521 U.S. at 874–79 (indecentcy provisions of the Communications Decency Act); *Playboy*, 529 U.S. at 813–27 (signal bleed provision); *Ashcroft v. ACLU*, 542 U.S. at 665–73 (Child Online Protection Act). The pattern is old. See *Butler v. Michigan*, 352 U.S. 380, 383 (1957) (a State may not “reduce the adult population . . . to reading only what is fit for children”).

120. Qian Li, Han Geng & Xin Hu, *Human-like Conversational Agents as Social Partners: A Scoping Review of Socioaffective Mechanisms, Well-Being Outcomes, Risks and Governance in the Post-Turing Era*, 9 *Frontiers Artificial Intelligence* (2026), <https://doi.org/10.3389/frai.2026.1810097> (PRISMA-ScR review of 58 sources identifying anthropomorphism, social presence, self-disclosure, parasocial attachment, and sycophancy as the operative mechanisms).

121. On sycophancy as a general property of models trained with reinforcement learning from human feedback, driven in part by preference judgments that systematically favor agreeable responses, see Mrinank Sharma et al., *Towards Understanding Sycophancy in Language Models*, arXiv:2310.13548 (2023). The finding matters for statutory design because a mechanism located in the training objective is not confined to the product class A2710 defines, and is not removed by prohibiting that class.

122. The mechanism literature is also more specific than the bill’s premise. A computational analysis of over thirty thousand shared conversations found emotional mirroring and synchrony resembling human bonding, with users skewing young, male, and toward maladaptive coping, and with some exchanges reproducing the dynamics of a harmful relationship. Minh Duc Chu et al., *Illusions of Intimacy: Emotional Attachment and Emerging Psychological Risks in Human-AI Relationships*, arXiv (2025). Two empirical studies locate the pathway in deficient self-regulation rather than in the product category, with the effect concentrated among lonely users. Lara Fröbel et al., *On the Dark Side of AI Companions*, in *Proceedings of the Annual Hawaii International Conference on System Sciences* (2026). And a relationship-science synthesis concludes that the operative research question is for whom these relationships are beneficial or harmful, which is a question about users and design rather than about the existence of the category. Molly Grace Smith, Thomas N. Bradbury & Benjamin R. Karney, *Can Generative AI Chatbots Emulate Human Connection? A Relationship Science Perspective*, 20 *Persp. on Psych. Sci.* 1081 (2025), <https://doi.org/10.1177/17456916251351306>.

D. O'Brien Does Not Reach a Content Defined Prohibition

O'Brien applies where the governmental interest is unrelated to the suppression of free expression.¹²³ A2710's interest is stated in its findings as the protection of a correct understanding of romantic love, and its prohibition is defined by the content of the interaction: romantic, intimate, or sexual.¹²⁴ A prohibition that must read the output to know whether it applies is content defined, whatever it is called.¹²⁵

The commercial speech route is worse for the state, not better. *Central Hudson's* intermediate standard still requires that the restriction be no more extensive than necessary,¹²⁶ and section 5 establishes that a less extensive restriction was available. The state cannot win the tailoring inquiry at any level of scrutiny while its own bill carries the narrower rule.

E. Severability Saves Section 5 and Concedes the Argument

The severability point is correct and unhelpful.¹²⁷ If section 4 falls and section 5 survives, New Jersey is left with a deception provision that would have reached the conduct in its own findings, which is the outcome this Commentary recommends. Severability is not a defense of the ban. It is a description of what remains when the ban is struck.

The narrowing construction argument fares no better.¹²⁸ Reading “designed or marketed to” as a purpose requirement confines the statute to products built as companions, which are precisely the products that operate age gates and crisis protocols, and leaves untouched the general models and community fine-tunes that operate neither. The narrowing construction makes the inversion worse.

123. *O'Brien*, 391 U.S. at 377.

124. Assemb. B. 2710 §§2(d), 3.

125. That is *Reed's* test, and *Paxton* reaffirmed it in the same breath as the incidental-burden carve-out. *Paxton*, 606 U.S. 461 (slip op. at 18) (citing *Reed v. Town of Gilbert*, 576 U.S. 155, 163–64 (2015)).

126. *Cent. Hudson*, 447 U.S. at 566.

127. Severability analysis in New Jersey asks whether the legislature would have enacted the remainder standing alone. Section 5 has an independent scienter element and an independent enforcement track through the Consumer Fraud Act, so it survives that inquiry comfortably. See Assemb. B. 2710 §§5–6.

128. See *supra* Part III (every enacted statute regulates precisely the products a purpose-limited reading would leave covered).

XI. The Statute New Jersey Could Have Written

A. Disclosure

A requirement that a system interacting directly with a person identify itself as artificial reaches the conduct in A2710's findings directly, because the harm there was a machine asserting that it was a person.¹²⁹ It is the approach New York, California, Washington, Oregon, and Connecticut adopted, and it is the approach the European Union codified as a transparency obligation on providers of systems intended to interact directly with natural persons.¹³⁰ A compelled factual disclosure of the speaker's non-human status is the kind of requirement that survives review most comfortably, because it adds information rather than removing expression.¹³¹

B. The Transaction

The one point in an open distribution chain where a regulator retains genuine visibility is payment.¹³² A gated download for consideration produces a named payee, a processor, a transaction record, and a receipt. Age attestation and disclosure attached to the transaction reach a population the artifact rule cannot, and regulate commercial conduct rather than expressive output. A2710 does nothing with the transaction. The distinction is defensible on its own terms and has been argued for independently: the insertion of undisclosed promotional content into an affective exchange collapses the boundary between a market transaction and a communicative

129. The Reuters reporting establishes both halves: the persona asserted personhood, and the internal standards governing it did not forbid the assertion. *See supra* note 1.

130. Regulation (EU) 2024/1689, art. 50, 2024 O.J. (L 1689) (transparency obligations for providers and deployers of certain AI systems); *see also* Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, 2026 O.J. (L 1081).

131. *Zauderer v. Off. of Disciplinary Couns.*, 471 U.S. 626, 651 (1985). The limit is real and should be respected in drafting: a disclosure mandate that is unduly burdensome or that compels contested rather than factual and uncontroversial content receives heightened scrutiny. *Nat'l Inst. of Fam. & Life Advoc. v. Becerra*, 585 U.S. 755, 768–73 (2018).

132. Payment is the chokepoint states already use where distribution is otherwise diffuse. *See supra* Part VI.D. It also has the advantage of reaching commercial conduct rather than expressive output, which places the regulation on the intermediate side of *Paxton*'s sort rather than the strict side. *See supra* Part VIII.

one, and a firm legal line between commercial and non-commercial conversational contexts is a precondition for stabilizing these products.¹³³

C. Design Duty

The most durable route attaches a duty of care to design choices rather than to expressive output, and a federal court has held the distinction workable in this exact setting.¹³⁴ *Garcia* held that a companion application is a product for product liability purposes insofar as the claims arise from defects in its design, such as age verification and reporting mechanisms, rather than from ideas or expressions within it.¹³⁵ The court further found a duty on the ground that releasing the system created a foreseeable zone of risk the defendants were positioned to control.¹³⁶

That route also survives Section 230, which does not bar a claim premised on a defendant's own design conduct rather than on its role as publisher of third party content.¹³⁷ A2710 selected the one framing that guarantees the intermediary is never reached and the design question is never asked.

The design-duty route is where this Commentary stops and where the regulatory-theory literature says the work begins. Fraser and colleagues classify the 2026 statutes into access gating, measures aimed at relationship features, and process accountability, conclude that an effective regime needs all three, and add that even all three leave the underlying condition untouched, because none of them addresses the asymmetry created when a provider controls

133. M. Figueroa, *The Limits of Artificial Companionship*, arXiv (2026).

134. The design-duty route is also the one the scholarship converges on from three different starting points. See Maitra, *supra* note 9 (design-and-disclosure statutes survive intermediate scrutiny while flat bans do not); Huang, *supra* note 7 (institutional autonomy conditional on a self-governance system that mitigates speech harms); Nash & Felton, *infra* note 77 (systems-based regulation of platform design outperforms content control).

135. *Garcia*, 785 F. Supp. 3d at 1195.

136. *Id.* at 1197–99 (applying *McCain v. Fla. Power Corp.*, 593 So. 2d 500, 502–03 (Fla. 1992)).

137. See *supra* Part IX.A (the state's coverage argument), and *supra* Part X.A (why winning it would not win the case). *Lemmon v. Snap, Inc.*, 995 F.3d 1085, 1092–94 (9th Cir. 2021); cf. *Zeran v. Am. Online, Inc.*, 129 F.3d 327, 330 (4th Cir. 1997) (stating the publisher immunity the design theory avoids). Treatment note: *Lemmon* carries a caution signal and the design conduct line is contested in application. It has been applied favorably, see *Est. of Bride v. Yolo Techs., Inc.*, 112 F.4th 1168 (9th Cir. 2024); *In re Soc. Media Adolescent Addiction/Personal Inj. Prods. Liab. Litig.*, 702 F. Supp. 3d 809 (N.D. Cal. 2023), and distinguished by state appellate courts, see *Patterson v. Meta Platforms, Inc.*, 2025 N.Y. Slip Op. 04438 (App. Div. 2025) (majority declining to sustain design claims that the dissent would have permitted under *Lemmon*). The route is available, not automatic.

intimacy at scale. Their proposal is a general, open-ended duty of care.¹³⁸ That proposal is broader than anything argued here, and this Commentary takes no position on it. It is worth naming because it marks the ceiling of the instrument discussion: the three mechanisms specified in this Part are the floor a state should not fall below, not the limit of what a state might do.

XII. Position Taking

There is a name for legislation whose return is delivered by its announcement rather than by its operation.¹³⁹ Political science calls it position taking, and the classic account holds that the electoral reward attaches to the public statement itself, so that a legislator is compensated as a speaker rather than as a doer.¹⁴⁰ The legal academic version studies the expressive function of law, the function of law in making statements as opposed to controlling behavior directly.¹⁴¹

Sunstein is not hostile to the category. He treats expressive legislation as a legitimate instrument for changing norms, and adds the qualification that governs here: legal statements that produce bad consequences should not be enacted even where they seem reasonable or noble.¹⁴² That qualification is the case against A2710.¹⁴³ The bill does not merely fail to accomplish its stated aim. It would clear the field of the providers a regulator can reach and leave the ones it cannot, with no age provision on either side of the line.

138. Fraser, Szczuka & Ciriello, *supra* note 9.

139. The point is descriptive rather than an imputation of motive. Nothing in this Commentary depends on the sponsor's intent, and the analysis would be identical if the bill were drafted in complete good faith, because the defects are structural.

140. DAVID R. MAYHEW, *CONGRESS: THE ELECTORAL CONNECTION* (Yale Univ. Press 1974). The position-taking discussion appears in the chapter on the electoral incentive, alongside advertising and credit claiming; the page range is left unpinned because the 1974 pagination could not be verified against any full-text index reachable in this build.

141. Cass R. Sunstein, *On the Expressive Function of Law*, 144 U. Pa. L. Rev. 2021 (1996). First page, volume, issue, and date confirmed against the Crossref record for DOI 10.2307/3312647; the article's interior pagination is not carried in any index reachable in this build, so citations to it here are general rather than pinned.

142. *Id.*

143. *See supra* Parts VI.C, VII.A.

Two structural features insulate the drafting from correction.¹⁴⁴ Legislators are absolutely immune from suit for legislative acts, and New Jersey's constitution provides that for any statement, speech, or debate in either house or at any meeting of a legislative committee, members shall not be questioned in any other place.¹⁴⁵ And the standard of review is forgiving in a way administrative law is not. An agency promulgating this rule would face arbitrary and capricious review, obliging it to examine the relevant data and articulate a rational connection between the facts found and the choice made.¹⁴⁶ A legislature enacting the identical rule need articulate no reason at all, and a reviewing court is obliged to hypothesize one on its behalf.¹⁴⁷

XIII. Conclusion

A2710 treats distribution as though it happens at a storefront, in a world where it happens at a repository, a tunnel, and a payment gate.¹⁴⁸ Everything else follows from that mismatch: four verbs that no longer sort anyone, enforcement that falls by collectability, an age gate that disappears along with the regulated products, and a narrower rule the drafter wrote into section 5 and then declined to rely on.¹⁴⁹

The underlying concern is serious and the evidentiary picture remains thin, which is a reason for careful instruments rather than for none.¹⁵⁰ Fourteen legislatures, a Congress, and a European regulator reached the same record within eighteen months and every one of them chose disclosure, crisis protocol, age assurance, or design duty. A statute keyed to disclosure at the start

144. Neither is a defect in the bill; both are features of the system the bill sits in, and they explain why the correction has to come from a court or from the sponsor rather than from review. *See infra* Part XI.

145. N.J. CONST. art. IV, §4, para. 9; *Tenney v. Brandhove*, 341 U.S. 367, 372–79 (1951); *see also Bogan v. Scott-Harris*, 523 U.S. 44, 49 (1998).

146. *Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co.*, 463 U.S. 29, 43 (1983).

147. *Beach Commc'ns*, 508 U.S. at 315.

148. *See supra* Part VI.

149. *See supra* Parts IV, V, VI.C, VII.

150. *See supra* Part X.C (why uncertainty about magnitude does not license direction). Thin does not mean absent. The first peer-reviewed clinical case report of new-onset AI-associated psychosis appeared in 2025, and clinicians have called publicly for systematic analysis of chat logs as forensic input. The absence of base rates is an argument for instruments that generate observability, which disclosure and reporting duties do and a prohibition does not.

of a conversation, to age assurance at the transaction, and to a duty of care over design choices would reach the death recorded in A2710's own findings, would survive constitutional review, and would not clear the field of the only providers a regulator can serve.

New Jersey could write that bill. Its neighbors already did.¹⁵¹

151. New York's statute took effect on November 5, 2025, three weeks before A2710's predecessor was introduced and two months before A2710 itself. *See supra* note 36. The model was available across the Hudson before the pen touched the page.

References

- Artificial Intelligence Companion Models, N.Y. Gen. Bus. Law §§1700–1708 (McKinney 2025).
- Ashcroft v. American Civil Liberties Union, 542 U.S. 656 (2004).
- Ashcroft v. Free Speech Coalition, 535 U.S. 234 (2002).
- Assembly Bill 2710, 222d Leg., Reg. Sess. (N.J. 2026) (“Authentic Relationships Act”).
- Bantam Books, Inc. v. Sullivan, 372 U.S. 58 (1963).
- Bassini, M. (2025). Speech without a speaker: Constitutional coverage for generative AI output? *European Constitutional Law Review*, 21, 1. <https://doi.org/10.1017/s1574019625100771>.
- Bernstein v. United States Department of Justice, 176 F.3d 1132 (9th Cir. 1999).
- Bogan v. Scott-Harris, 523 U.S. 44 (1998).
- Brown v. Entertainment Merchants Association, 564 U.S. 786 (2011).
- Butler v. Michigan, 352 U.S. 380 (1957).
- Center for Democracy & Technology. (2026). *2026 State and Federal AI Legislation Updates*. <https://cdt.org/insights>.
- Central Hudson Gas & Electric Corp. v. Public Service Commission, 447 U.S. 557 (1980).
- Chu, M. D., et al. (2025). *Illusions of intimacy: Emotional attachment and emerging psychological risks in human-AI relationships*. arXiv.
- City of Chicago v. Morales, 527 U.S. 41 (1999).
- Companion Chatbots, Cal. S.B. 243, 2025–2026 Reg. Sess. (Cal. 2025).
- Craig v. Boren, 429 U.S. 190 (1976).
- Crime and Policing Act 2026 (U.K.).
- D’Agostino v. Maldonado, 216 N.J. 168, 78 A.3d 527 (2013).
- Eisenstadt v. Baird, 405 U.S. 438 (1972).
- Eom, D., Renner, J., & Chinn, S. (2026). *Intimacy as service, harm as externality: Critical perspectives on AI companion platform accountability*. arXiv. <https://doi.org/10.48550/arxiv.2604.06381>.
- Estate of Bride v. Yolo Technologies, Inc., 112 F.4th 1168 (9th Cir. 2024).
- Federal Communications Commission v. Beach Communications, Inc., 508 U.S. 307 (1993).
- Figueroa, M. (2026). *The limits of artificial companionship*. arXiv.
- Fisher, S. A., Howard, J. W., & Kira, B. (2024). Moderating synthetic content: The challenge of generative AI. *Philosophy & Technology*, 37, 1. <https://doi.org/10.1007/s13347-024-00818-9>.
- Fröbel, L., et al. (2026). On the dark side of AI companions: How a parasocial preference for a social chatbot can lead to pathological chatbot use. In *Proceedings of the Annual Hawaii International Conference on System Sciences*.
- Fraser, H., Szczuka, J. M., & Ciriello, R. (2026). Regulating artificial intimacy: From locks and blocks to relational accountability. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1476–1493). <https://doi.org/10.1145/3805689.3806790>.

- Free Speech Coalition, Inc. v. Paxton, 606 U.S. 461 (2025).
- Future of Privacy Forum. (2026). *The Chatbot Moment: Mapping the Emerging 2026 U.S. Chatbot Legislative Landscape*. <https://fpf.org/blog>.
- Garcia v. Character Technologies, Inc., 785 F. Supp. 3d 1157 (M.D. Fla. 2025).
- Ginsberg v. New York, 390 U.S. 629 (1968).
- Grayned v. City of Rockford, 408 U.S. 104 (1972).
- Guidelines for User Age-verification and Responsible Dialogue Act of 2026, S. 3062, 119th Cong. (2026).
- Home Office. (2026). *Crime and Policing Act 2026: Child Sexual Abuse Material Factsheet*. <https://www.gov.uk/government/publications>.
- Horwitz, J. (2025). Meta's flirty AI chatbot invited a retiree to New York. He never made it home. *Reuters*. <https://www.reuters.com/technology>.
- Huang, T. (2026). *Generative AI as an institutional tool under the First Amendment*. <https://doi.org/10.2139/ssrn.6311158>.
- In re Social Media Adolescent Addiction/Personal Injury Products Liability Litigation, 702 F. Supp. 3d 809 (N.D. Cal. 2023).
- Junger v. Daley, 209 F.3d 481 (6th Cir. 2000).
- Kolender v. Lawson, 461 U.S. 352 (1983).
- Lamont v. Postmaster General, 381 U.S. 301 (1965).
- Lemmon v. Snap, Inc., 995 F.3d 1085 (9th Cir. 2021).
- Li, Q., Geng, H., & Hu, X. (2026). Human-like conversational agents as social partners: A scoping review of socioaffective mechanisms, well-being outcomes, risks and governance in the post-Turing era. *Frontiers in Artificial Intelligence*, 9. <https://doi.org/10.3389/frai.2026.1810097>.
- Maitra, S. (2026). *Brown's children: AI chatbots, minors, and the First Amendment*. <https://doi.org/10.2139/ssrn.6915498>.
- Martin v. City of Struthers, 319 U.S. 141 (1943).
- Mayhew, D. R. (1974). *Congress: The Electoral Connection*. Yale University Press.
- McCain v. Florida Power Corp., 593 So. 2d 500 (Fla. 1992).
- Mei, Y. (2026). *The disintegration of free speech*. <https://doi.org/10.2139/ssrn.6243039>.
- Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd., 545 U.S. 913 (2005).
- Moody v. NetChoice, LLC, 603 U.S. 707 (2024).
- Morgan Lewis. (2026). *Washington and Oregon Regulate AI Companions: Key Compliance Changes*. <https://www.morganlewis.com/pubs>.
- Morrison & Foerster. (2025). *New York and California Enact Landmark AI Companion Laws: What Operators Need to Know*. <https://www.mofo.com/resources/insights>.
- Motor Vehicle Manufacturers Association v. State Farm Mutual Automobile Insurance Co., 463 U.S. 29 (1983).
- Nash, V., & Felton, L. (2024). Treating the symptoms or the disease? Analysing the UK Online Safety Act's approach to digital regulation. *Policy & Internet*, 16, 818.
- National Institute of Family and Life Advocates v. Becerra, 585 U.S. 755 (2018).
- National Pork Producers Council v. Ross, 598 U.S. 356 (2023).

- NetChoice v. Fitch, 145 S. Ct. 2658 (2025).
- New Jersey Constitution art. IV, §4, para. 9.
- New Jersey Consumer Fraud Act, N.J. Stat. Ann. §§56:8-1 to -227.
- New York v. Ferber, 458 U.S. 747 (1982).
- Online Safety and Artificial Intelligence Regulation Act, Conn. Pub. Act No. 26-15 (2026).
- Orrick, Herrington & Sutcliffe LLP. (2026). *2026 State Chatbot Laws: Key Provisions and Regulatory Trends*. <https://www.orrick.com/en/Insights>.
- Papachristou v. City of Jacksonville, 405 U.S. 156 (1972).
- Patterson v. Meta Platforms, Inc., 2025 N.Y. Slip Op. 04438 (App. Div. 2025).
- Prince v. Massachusetts, 321 U.S. 158 (1944).
- Reed v. Town of Gilbert, 576 U.S. 155 (2015).
- Regulating Artificial Intelligence Companion Chatbots, Wash. H.B. 2225, 69th Leg., Reg. Sess. (Wash. 2026).
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), 2024 O.J. (L 1689).
- Relating to Artificial Intelligence Companions, Or. S.B. 1546, 83d Leg. Assemb., Reg. Sess. (Or. 2026).
- Reno v. American Civil Liberties Union, 521 U.S. 844 (1997).
- Sable Communications of California, Inc. v. Federal Communications Commission, 492 U.S. 115 (1989).
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askill, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2023). Towards understanding sycophancy in language models. *arXiv*. <https://doi.org/10.48550/arXiv.2310.13548>.
- Smith, M. G., Bradbury, T. N., & Karney, B. R. (2025). Can generative AI chatbots emulate human connection? A relationship science perspective. *Perspectives on Psychological Science*, 20(6), 1081–1099. <https://doi.org/10.1177/17456916251351306>.
- Stanley v. Georgia, 394 U.S. 557 (1969).
- Sunstein, C. R. (1996). On the expressive function of law. *University of Pennsylvania Law Review*, 144(5), 2021. <https://doi.org/10.2307/3312647>.
- Tenney v. Brandhove, 341 U.S. 367 (1951).
- Thiedemann v. Mercedes-Benz USA, L.L.C., 183 N.J. 234, 872 A.2d 783 (2005).
- Thurman, N., Nalmpatian, A., & Obster, F. (2022). Lessons from France on the regulation of Internet pornography: How displacement effects, circumvention, and legislative scope may limit the efficacy of Article 23. *Policy & Internet*, 14, 690. <https://doi.org/10.1002/poi3.293>.
- Transparency Coalition. (2026). *Watershed Year for Chatbot Safety: 14 New State Laws Passed So Far in 2026*. <https://www.transparencycoalition.ai/news>.
- Troutman Amin. (2026). *Analyzing the New AI Companion Chatbot Laws*. <https://www.troutmanprivacy.com>.
- Turner Broadcasting System, Inc. v. Federal Communications Commission, 512 U.S. 622 (1994).
- Turner Broadcasting System, Inc. v. Federal Communications Commission, 520 U.S. 180 (1997).
- United States v. O'Brien, 391 U.S. 367 (1968).

United States v. Playboy Entertainment Group, Inc., 529 U.S. 803 (2000).
United States v. Stevens, 559 U.S. 460 (2010).
United States v. Williams, 553 U.S. 285 (2008).
Unsloth. (2026). *Introducing Unsloth Desktop*. <https://unsloth.ai/docs/desktop>.
Vermont Act 156, H.816, 2025–2026 Leg. Sess. (Vt. 2026).
Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council, Inc., 425 U.S. 748 (1976).
Weidmann, P. E. (2026). *Heretic: Fully automatic censorship removal for language models*. GitHub.
<https://github.com/p-e-w/heretic>.
Weinberg v. Sprint Corp., 173 N.J. 233, 801 A.2d 281 (2002).
West Virginia State Board of Education v. Barnette, 319 U.S. 624 (1943).
Zauderer v. Office of Disciplinary Counsel of the Supreme Court of Ohio, 471 U.S. 626 (1985).
Zeran v. America Online, Inc., 129 F.3d 327 (4th Cir. 1997).

☛ ☚