

Safety Theater:

The People-First Chatbot Act and the Protection It Withholds

Travis Gilly*. July 2026

A Commentary on H.R. 9619, the People-First Chatbot Act, 119th Congress, as introduced.

Abstract

On July 9, 2026, Representatives Valerie Foushee and Greg Casar introduced H.R. 9619, the People-First Chatbot Act, federal legislation built on model text drafted by the Consumer Federation of America, the Electronic Privacy Information Center, and Fairplay. The bill is promoted as a comprehensive answer to chatbot harm: data privacy, safety by design, strict products liability, and a private right of action for injured users.

Read at the provision level, the statute delivers something different. It regulates a conversational surface rather than a technology, leaving every consequential deployment of the same models untouched. Its privacy sections bar advertisers and the government from chat logs while its safety section permits those logs to be retained for as long as five years and requires a monthly risk assessment whose metrics the Federal Trade Commission will define only after passage. Its liability regime rests on a metric deferred to that Commission and an “injury in fact” the text never scopes, and imposes liability without regard to care on systems whose outputs cannot be guaranteed. Its remedy reaches no bystander: the injury liability runs to a user harmed through use, and the private right of action reaches only a person injured by one of the Act’s own violations, so the gravest documented victims of chatbot involved violence, the bystanders of the Florida State University and Tumbler Ridge shootings, take nothing, though ordinary products liability doctrine already serves them. It regulates the wrong actor as well as the wrong unit: every duty attaches to a provider who creates or distributes a chatbot, the word “deployer” appears nowhere, and a person running an autonomous agent on his own account under his own name owes nothing while the open-source maintainer answers with reasonable care no defense and privity expressly abolished. And the consent on which its entire adult privacy title depends cannot be obtained from part of the class it names: the Act makes “informed” a condition of consent, elects a written description in easy-to-understand language as the vehicle, and imports the Americans with Disabilities Act’s definition of disability, which is in substantial part a definition of substantial limitation in reading and thinking, while the peer-reviewed record on simplified consent information reports null results for the readers who need it most. And if the statute’s “injury in fact” reaches emotional harm, chatbot users would recover the emotional distress damages that *Cummings v. Premier Rehab Keller* categorically denies to disabled plaintiffs suing under Section 504.

*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.

This Commentary supplies a provision-level critique of the introduced text of H.R. 9619, documents the bystander gap against the litigation now pending over the Florida State and Tumbler Ridge attacks, identifies the emotional distress inversion the bill would work against disability civil rights remedies, and closes with a seven-part alternative, each element with a working model in enacted law: regulate by consequence, minimize data rather than retain and scan it, keep safety detection in the moment, place the substance in the statute, match liability to the machine, make the remedy usable, and invest in the capability of the protected.

Keywords: People-First Chatbot Act; H.R. 9619; chatbot regulation; affirmative consent; cognitive accessibility; delay aversion; deployer liability; agentic AI; strict products liability; bystander liability; private right of action; data retention; safety theater; AI-induced psychosis; AI literacy; *Cummings v. Premier Rehab Keller*; *Elmore v. American Motors Corp.*

Contents

I	Introduction	4
II	The Wrong Unit of Regulation	5
A	The Interface, Not the Model	5
B	The Provider, Not the Deployer	11
III	The Privacy Inversion.	14
A	What the Privacy Sections Provide	14
B	What the Safety Section Builds.	15
C	The Deferred Target.	16
D	The Sign on the Door.	17
E	The Consent That Cannot Be Given	19
1	What the Record Says About Easy-to-Understand Language	20
2	Position, Not Content	22
3	The Inference the Instrument Requires	24
4	The Answer They Drafted and Spent on Children	25
5	Two Standards and No Limit	26
6	What the Common Law Already Does, and What the Act Does Beyond It	27
7	What Follows	29
IV	Two Deferred Terms: The Assessment Metrics and Injury	30
A	The Metrics, Defined Later.	30
B	Injury, Left Unspecified	30
V	A Liability Regime Built for a Different Machine	31
VI	Remedies That Do Not Reach the Injured.	32
A	A Remedy Few Can Lift	32
B	The Emotional Distress Inversion	33
C	The Bystander Gap	35
D	Redundant with Better Law	37
VII	Marketing a Condition the Statute Does Not Name	38
VIII	Restriction Without Capability.	40
IX	A Better Path	40
A	Regulate by Consequence Rather than Format	41
B	Minimize Data Rather than Retain and Scan It.	42
C	Keep Safety Detection in the Moment.	43
D	Put the Substance in the Statute	44
E	Match Liability to the Machine.	44
F	Make the Remedy Usable.	45
G	Invest in the Capability of the Protected	45
X	Conclusion	46

I. Introduction

On July 9, 2026, Representatives Valerie Foushee and Greg Casar introduced H.R. 9619, the People-First Chatbot Act, in the House of Representatives.¹ The sponsors describe the bill as comprehensive legislation establishing federal safeguards for Americans who interact with artificial intelligence chatbots: control over personal data, disclosure of machine identity, protections for children and vulnerable users, and meaningful recourse when chatbots cause harm.² The bill is built on model legislation released in January 2026 by the Consumer Federation of America, the Electronic Privacy Information Center, and Fairplay, and already introduced in several state legislatures.³ The advocacy coalition behind the model presents the federal bill as the legislation that “kids and adults urgently need,” naming among the harms it answers “suicide, AI psychosis, compulsive use, and unhealthy emotional dependency.”⁴

The goal is right and overdue. The statute as drafted does not meet it. This Commentary examines the bill provision by provision and finds a measure that regulates the wrong unit, contradicts its own privacy promises, rests its liability regime on terms it never defines, offers

1. H.R. 9619, 119th Cong. (2026).

2. Press Release, Off. of Rep. Valerie Foushee, Reps. Foushee, Casar Introduce Legislation to Protect Children and Americans’ Privacy from AI Chatbot Harms and Require Chatbot Safety Assessments (July 9, 2026), <https://foushee.house.gov/media/press-releases/rep-foushee-casar-introduce-legislation-to-protect-children-and-americans-privacy-from-ai-chatbot-harms> [hereinafter Foushee Press Release].

3. CONSUMER FED’N OF AM., ELEC. PRIVACY INFO. CTR. & FAIRPLAY, PEOPLE-FIRST CHATBOT BILL (Jan. 20, 2026), <https://fairplayforkids.org/wp-content/uploads/2026/01/People-First-Chatbot-Bill-Jan-20-2026-1.pdf> (last visited July 11, 2026) [hereinafter Model Bill]; see also Grace Tepper, *Introducing: A People-First Chatbot Bill*, Benton Inst. for Broadband & Soc’y (Jan. 23, 2026), <https://www.benton.org/blog/introducing-people-first-chatbot-bill> (reproducing the model bill’s operative sections). This Commentary analyzes the introduced House text of H.R. 9619 as released by the sponsor’s office, https://foushee.house.gov/imo/media/doc/final_people-first_chatbot_act.pdf (last visited July 13, 2026), and cites its sections throughout; the model bill is cited where it supplies provenance or differs from the introduced text. All section references are to H.R. 9619 as introduced unless otherwise noted.

4. Press Release, Fairplay, People-First Chatbot Act Introduced in US House (July 10, 2026), <https://fairplayforkids.org/people-first-chatbot-act-introduced-in-us-house/> [hereinafter Fairplay Press Release]; see also Press Release, Consumer Fed’n of Am., CFA Celebrates House Introduction of the People-First Chatbot Act (July 10, 2026), <https://consumerfed.org/news/press-releases/cfa-celebrates-house-introduction-of-the-people-first-chatbot-act/> (describing the sponsors and supporting coalition).

a remedy the injured cannot use and the gravest victims cannot invoke, and markets a clinical condition its text does not name.⁵

The failures are not incidental. Each traces to the same drafting choice: the bill regulates the conversational surface where harm is most visible rather than the technology and the consequences where harm is most severe, and it defers the substance of its safety promise to a rulemaking that has not occurred. The result is legislation that performs protection. Part II examines the coverage choice. Part III examines the collision between the bill's privacy sections and the retention and scanning apparatus its safety section builds. Part IV examines the two deferred or undefined terms on which the operative provisions rest. Part V examines the strict liability regime against the behavior of the systems it would govern. Part VI examines who the remedy reaches and who it leaves out, measured against the litigation now pending over two mass shootings. Part VII examines the gap between the harm the coalition markets and the harm the text addresses. Part VIII examines what the bill declines to build. Part IX proposes the alternative.

II. The Wrong Unit of Regulation

A. The Interface, Not the Model

The bill's obligations attach to chatbots. The introduced text defines an artificial intelligence chatbot as any interactive computer service or software application that "generates responses that are not fully predetermined" and "accepts open-ended natural-language or multimodal user input and produces adaptive or context-responsive output," and it expressly excludes a system "the responses of which are limited to contextualized replies" that is "unable to respond on a range of topics outside of a narrow specified purpose."⁶ A separate inclusion reaches companion

5. See *infra* Parts II–VIII.

6. H.R. 9619, § 5(3)(A). This definition is materially narrower than the model bill's, which reached "any artificial intelligence, algorithmic, or automated system" that "simulates interpersonal interactions or conversation." See Model Bill, *supra* note 3. The change matters: a system whose responses are fully predetermined, and a narrow single-purpose tool, are carved out of the introduced text where the model text swept them in.

systems: those that simulate “a sustained interpersonal relationship or emotional interaction” through persistent affection or attachment, emotional disclosures from the user, or “at least one persistent identity, persona, or character.”⁷ The definition reaches beyond text to multimodal output that meets those criteria. What it does not reach is everything the same technology does without conversing.

The models that power consumer chatbots also screen job applicants, rank admissions candidates, score eligibility for public benefits, and inform government risk assessments. None of those systems simulates interpersonal interaction. Each returns a determination rather than a conversation. Under the bill’s definition, none is covered, though the underlying model may be identical and the decisions those systems produce alter employment, education, housing, and liberty. The pattern is not unique to this bill; the federal chatbot measures introduced across the 119th Congress share the same focus on the conversational surface.⁸ The bill therefore places its heaviest obligations, including strict liability, on the most visible deployment of the technology and leaves the consequential decision systems untouched.

The same criteria that let the decision systems escape sweep in tools that no one would call the target of a chatbot statute. The definition reaches any system that produces adaptive, non-predetermined output across a range of topics from open-ended input, and the bill exempts no category of use. A developer’s assistant such as Claude Code at the command line, the Codex command-line interface, or a coding assistant built into an editor such as Cursor, Zed, or Google’s Antigravity produces exactly that kind of output to do work that is not conversation, and each therefore satisfies the definition. The bill defines a chatbot provider as any person who makes a chatbot available and a user as any natural person regardless of age, and it contains

7. H.R. 9619, § 5(3)(B).

8. *See also* Curbing Harmful AI Tools By Offering Transparency Act, H.R. 7985, 119th Cong. (2026) (restricting the impersonation of licensed professionals in chatbot output and marketing); Children’s Health, Advancement, Trust, Boundaries, and Oversight in Technology Act, S. 4407, 119th Cong. (2026) (requiring family accounts and verifiable parental consent for minors’ chatbot use).

no applicability section and no exemption for professional, enterprise, or developer tools.⁹ On the text, a coding assistant owes its user the notice that it is a chatbot before every output and once an hour thereafter,¹⁰ must be assessed each month against harm metrics the rulemaker has not written,¹¹ may have its session logs retained for as long as five years,¹² and exposes its maker to liability for any injury in fact caused through use notwithstanding all reasonable care.¹³ The reach does not stop at software a developer runs. Because the definition extends to multimodal output, it draws in the general-purpose voice assistant shipped on every phone: Apple would become a chatbot provider because Siri answers open-ended questions across a range of topics in a human voice, carrying the same hourly disclosure, five-year retention, monthly assessment, and no-fault liability. The introduced text does carve out the narrow single-purpose tool, the reminder-only or balance-inquiry bot “unable to respond on a range of topics outside of a narrow specified purpose,”¹⁴ but it leaves every general assistant inside. Either the drafters intend that result, which is difficult to credit, or they wrote a definition they did not follow to its consequences. The coverage is therefore overbroad and underinclusive at once, the signature of a unit chosen for where harm is visible rather than for where it occurs.

The catalog does not end with tools a person runs on a device, though here the introduced text draws a line the model text did not. A narrow customer-service bot that only answers a fixed band of questions before passing a caller to a live agent may fall within the carve-out for systems “unable to respond on a range of topics outside of a narrow specified purpose.”¹⁵ But the general-purpose assistant now embedded across retail, banking, and government sites, the one that answers open-ended questions across subjects, is not narrow, and it is covered. A person

9. H.R. 9619, §§ 5(3), (4), (22) (defining “artificial intelligence chatbot,” “artificial intelligence chatbot provider,” and “user,” the last as an individual regardless of age, and containing no applicability limitation or exemption for professional, enterprise, or developer tools).

10. *Id.* § 2(b)(2).

11. *Id.* §§ 2(c)(1), 3(a)(3).

12. *Id.* § 2(a)(1)(F).

13. *Id.* § 4(d).

14. *Id.* § 5(3)(A)(ii).

15. H.R. 9619, § 5(3)(A)(ii).

who asks a benefits or tax question of a federal agency's general assistant is a user under the Act, and that assistant now very often runs on a commercial model the agency licenses through the General Services Administration, which places the chatbot provider squarely within reach as a company rather than as a government body.¹⁶ The person's exchange with a government service then becomes a record that may be kept for as long as five years and is subject to the monthly assessment. Enterprise assistants reach the same way, so a workplace tool such as Notion, which answers plain-language questions about a team's own documents and beyond, is covered.¹⁷ So are the research assistants that lawyers now use, among them Thomson Reuters' CoCounsel, where an attorney describes a matter in ordinary language and works out litigation strategy, deposition questions, and document review through the exchange.¹⁸ On the text, the work product built in that exchange becomes a record that may be kept for as long as five years and is subject to the monthly assessment. The bill exempts nothing for business, enterprise, or professional use, and none of these general tools is beyond a consumer's reach, since a solo practitioner can subscribe to CoCounsel as readily as anyone buys the rest.

Even for the systems it means to reach, the bill assumes a human who is often not there. Its central duties run to a user, a natural person, and require telling that person before every output and once an hour that a chatbot is speaking.¹⁹ A growing share of use no longer fits that picture. A person can drive the same subscription through a headless command-line tool inside an automated workflow, where a hook or a scheduled agent, and not a human typing in real time, issues each prompt, and no human need read the output as it returns. The provider cannot tell the two apart. The same model, returning the same text, meets a live person in one session and a script in the next, with nothing to separate them, so the hourly notice reaches no one and the idea

16. See U.S. Gen. Servs. Admin., *GSA Propels Government into AI Revolution with Addition of Leading Solutions to Multiple Award Schedules* (Aug. 5, 2025) (adding Anthropic's Claude, Google's Gemini, and OpenAI's ChatGPT to the Multiple Award Schedule for use across federal agencies); GSA's OneGov agreements offered these models to agencies for nominal annual fees. The agency has stated the aim is to reshape "citizen experiences" and service delivery.

17. See Notion, *Chat with Notion AI*, <https://www.notion.com/product/ai> (last visited July 11, 2026).

18. See *CoCounsel Legal*, Thomson Reuters, <https://legal.thomsonreuters.com/en/products/cocounsel-legal> (last visited July 11, 2026).

19. See H.R. 9619, §§ 2(b)(2), 5(22).

of a user comes apart. Operative commands that presume a reader who is frequently absent were not written for the technology they govern.

The overbreadth is not academic, because of who relies on the tools it reaches. The command line and the integrated development environment are the entry point to computer science, software development, and engineering, and to nearly any technical path that touches code. The people who most need to grow fluent in them are adolescents at the age when a person first chooses a direction for working life. This is not a claim that young children belong at a terminal. It is a claim about teenagers deciding whether these fields are for them, for whom a regime that brands the coding assistant a chatbot to be labeled every hour, logged for five years, and litigated over teaches dependence in place of capability, and does so to the very students a protective statute should be equipping. Part VIII returns to the capability the bill declines to build.²⁰

A drafter might reply that no court would apply the definition so broadly, that a sensible judge would read it down to the companion systems the bill is really aimed at. That reply misjudges the present bench. The dominant method of statutory interpretation is textualism, under which a court enforces the words Congress enacted, an express statutory definition above all, even where the application would have surprised the drafters.²¹ The Supreme Court did precisely that in *Bostock v. Clayton County*, holding that the plain terms of Title VII governed an application its drafters did not foresee, because “the limits of the drafters’ imagination supply no reason to ignore the law’s demands.”²² Which method a given court will bring to a given statute a drafter cannot predict, and that uncertainty is itself the argument for precision. Justice Kagan, who pronounced textualism ascendant, has since written that the Court is textualist

20. *See infra* Part VIII.

21. The Antonin Scalia Lecture: A Dialogue with Justice Elena Kagan on the Reading of Statutes, Harvard Law School (Nov. 25, 2015) (Justice Kagan remarking that “we’re all textualists now”).

22. *Bostock v. Clayton County*, 590 U.S. 644, 660 (2020).

only when that suits it, and the interpretive debate she named remains unsettled.²³ A statute that depends on judges declining to enforce its own words is built on hope, and a drafter cannot choose which reader the statute will meet, so the words must be written for both, the textualist and the purposivist alike. A definition drawn with that care yields a single meaning, because it says what it means, and that is the standard the alternative in Part IX is built to meet.²⁴

The coverage choice fails even inside its own category. Consider a hiring manager who pastes applications into a general-purpose chatbot and asks which candidates to reject, and the chatbot's answer is discriminatory. The harm ran through a covered chatbot. Yet the bill's injury liability runs to "a user" harmed "through the use" of the chatbot, and its private right of action reaches only "a person injured by an act or practice in violation" of the Act.²⁵ The user in that scenario is the hiring manager, who suffered no injury. The rejected applicant suffered the injury, but he never used the chatbot, so the injury-liability track does not reach him, and the discriminatory output is not itself one of the Act's enumerated violations, so the private-right-of-action track does not reach him either. He holds no claim under the bill. The same structure repeats wherever a person in authority uses a chatbot to decide something about someone else: a landlord screening tenants, a caseworker triaging benefits, a teacher grading students. The interaction does not have to happen to the person the harm reaches, and when it does not, this bill's protection evaporates. Part VI returns to this gap where the stakes are highest.²⁶

23. *West Virginia v. EPA*, 597 U.S. 697 (2022) (Kagan, J., dissenting) ("It seems I was wrong. The current Court is textualist only when being so suits it."); *see also Are We All Purposivists Now?*, Harvard Law School (Apr. 3, 2025) (reporting Judge Rachel Kovner's 2025 Scalia Lecture, which argues that purposivism coexists with textualism on the Court today).

24. *See infra* Part IX (regulating by consequence and defining the operative terms in the statute itself).

25. H.R. 9619, §§ 4(c)(1) (private right of action for "a person injured by an act or practice in violation of section 2"), 4(d) (injury liability for "a user . . . who suffers an injury in fact caused through the use").

26. *See infra* Part VI.C.

B. The Provider, Not the Deployer

The unit problem has a second half. Having chosen the interface as its object, the bill then chooses supply as the only relationship that creates a duty, and every operative obligation in Section 2 attaches to an “artificial intelligence chatbot provider,” defined as “any person who creates, distributes (including to a third party), or otherwise makes publicly available an artificial intelligence chatbot.”²⁷ Three verbs, all of them describing the act of supplying a chatbot to other people. None of them describes the act of aiming one at somebody. The word “deployer” does not appear in the bill.

The omission was a choice rather than an oversight, because the drafting vocabulary existed and was already before the same Congress. The CHATBOT Act, introduced in the House four months earlier, defines a “covered entity” as “an individual or company that deploys an AI chatbot,” and hangs its duties there.²⁸ H.R. 9619 reaches a deployer exactly once, in a subsection confined to business entities using a chatbot for customer service communications, which must disclose the machine at the start of the exchange and offer transfer to a human on request.²⁹ The drafters knew how to write a duty that runs to the party operating the thing. They wrote one, and scoped it to the help desk.

Everything outside the help desk falls through. Consider the ordinary case now arriving: a person runs an autonomous agent on his own logged-in social media account, and the agent posts, replies, and answers direct messages under his name. The software is squarely a chatbot under the Act, which sweeps in any covered application that “presents at least one persistent identity, persona, or character” in a manner simulating a sustained interpersonal relationship.³⁰ But the account holder supplies nothing. No third party can access his agent; they can only be addressed by it. He has not created it, does not distribute it, and does not make it publicly available, so he is

27. H.R. 9619, § 5(4).

28. CHATBOT Act, H.R. 7985, 119th Cong. § 5(5) (2026).

29. H.R. 9619, § 2(d).

30. *Id.* § 5(3)(B)(ii). The threshold is one persistent persona; an agent operating a named account satisfies it on the account name alone.

not a provider, no duty attaches to him, and no plaintiff has an action against him under this Act. The only party in the transaction who decided to point a machine at another human being is the only party the statute does not reach.

This is not a hypothetical technology. An open-source personal agent released in 2025 became one of the most-starred repositories on GitHub within months; it is self-hosted, keeps persistent memory, runs a daemon that acts without being prompted, and ships integrations for more than twenty messaging platforms.³¹ By February 2026 press accounts described such agents creating dating-app accounts on their owners' behalf without being asked.³²

The consequence is worse than a gap, because the Act does supply a defendant. It is the wrong one. The bill defines “user” as “an individual, regardless of age,”³³ which is not a definition but a synonym, so the stranger the agent messages is a user of it, notwithstanding that she opened no application, agreed to nothing, and has no relationship with anyone in the chain. The notice duty runs to her: a provider “shall provide a clear, conspicuous, and explicit notice to a user that the user is interacting with an artificial intelligence chatbot rather than a human” before the first output and every hour after.³⁴ No notice was given, so there is a violation, and the private right of action reaches “a person injured by an act or practice in violation” of Section 2.³⁵ The provider she must sue is whoever created the agent software and made it publicly available, which for an open-source project is its maintainer. The liability title then closes the trap: the provider answers for a user's injury in fact “even if” it “exercised all reasonable care in the design and distribution” and “even if” it “did not directly distribute the artificial intelligence chatbot to

31. See Peter Steinberger, *OpenClaw*, GitHub, <https://github.com/steipete/openclaw> (last visited July 14, 2026) (documenting the self-hosted architecture, the persistent-memory store, the unprompted heartbeat loop, and the messaging-platform plugins).

32. See *AI Agents Are Signing Their Owners Up for Dating Apps*, Taipei Times (Feb. 14, 2026), <https://www.taipeitimes.com/> (last visited July 14, 2026) (reporting unprompted account creation by consumer agent software); see also Straits Times (Feb. 14, 2026) (same reporting).

33. H.R. 9619, § 5(22). The definition supplies no limiting principle; it is a synonym for a natural person, and every duty keyed to the word therefore runs to every human being a covered chatbot addresses.

34. *Id.* § 2(b)(2).

35. *Id.* § 4(c)(1).

the user or otherwise enter into a contractual relationship with the user.”³⁶ Care is not a defense and privity is expressly abolished. An unpaid volunteer owes an hourly disclosure to a stranger he has never met, arising from conduct by an operator he cannot observe, and nothing he does in the design of his software discharges it. The person who chose to run the agent owes nothing.

The notice cannot be delivered in any event, because the provision that specifies it addresses a surface the duty-holder does not control. The disclosure must appear “in a font size easily readable by an average user, and not smaller than the largest font size of other text appearing on the interface on which the artificial intelligence chatbot is provided.”³⁷ Where the agent runs inside a social platform, that interface belongs to the platform and the account holder. The provider cannot set a font size on it, cannot insert a banner into it, and cannot condition the account holder’s use of it. The parties who could do all three hold no duty. This is not a drafting gap that a court will paper over with a reasonableness gloss; it is an obligation that the only party charged with it has no lawful means to perform.

State law drawn eight years earlier reaches further and still does not close the case. California’s bot-disclosure statute makes it unlawful for “any person” to use a bot to communicate with another person in California with intent to mislead about its artificial identity, which is a deployer duty in substance.³⁸ But it bites only where the purpose is to incentivize a commercial transaction or influence a vote in an election,³⁹ and its definition of “bot” reaches only an account whose activity is “not the result of a person” in substantially all respects.⁴⁰ The hybrid account, where a human writes some messages and an agent writes the rest, falls outside that definition, and falls outside H.R. 9619 for want of a deployer duty at all. A 2018 state law about election bots addresses the operator of a conversational machine more directly than a 2026 federal bill written to address conversational machines.

36. *Id.* § 4(d)(1)–(2).

37. *Id.* § 2(b)(2)(A).

38. Cal. Bus. & Prof. Code § 17941(a) (2026).

39. *Id.*

40. *Id.* § 17940(a).

The damages schedule confirms how little the bill's own drafters believe the disclosure title matters. A plaintiff may recover up to \$10,000 in total for all violations of Section 2(b), the entire transparency title, for everything.⁴¹ Violations of the minors design-safety subsection carry statutory minimums of \$50,000, rising to \$100,000 where the result is emotional dependence or compulsive usage and \$250,000 where the result is a covered harm.⁴² Disclosure is the promise in the press release and the cheapest, aggregate-capped violation in the statute.

III. The Privacy Inversion

A. What the Privacy Sections Provide

The bill's privacy sections are its strongest and most concrete. A provider may not process a user's chat log to select, target, or customize advertising.⁴³ A provider may not sell chat logs, may not use a minor's chat log or personal data for training, and may not train on an adult's data without affirmative consent.⁴⁴ A government entity may not compel production of or access to input data or chat logs "except as provided by a warrant issued by a court under section 2518 of title 18," the federal wiretap-warrant provision.⁴⁵ Those provisions are real, and they explain why the privacy community greeted the bill warmly.⁴⁶

Whether the wiretap warrant limit holds in practice is a separate question, and the recent record is not reassuring. Across 2025 and 2026, the Department of Homeland Security issued administrative subpoenas, self-signed demands that no judge reviews and that compel nothing until a court orders compliance, to obtain user data from the same companies this bill would

41. H.R. 9619, § 4(c)(1)(B)(ii).

42. *Id.* § 4(c)(1)(B)(iii)(I)–(III).

43. H.R. 9619, § 2(a)(1)(B) (barring processing of a chat log to "determine whether to display an advertisement," to select "a product, service, or category of product or service to advertise," or to "customize an advertisement").

44. H.R. 9619, §§ 2(a)(1)(E), (C) (prohibiting sale of chat logs, prohibiting training uses of the data of users known to be under eighteen, and conditioning training uses of adults' data on affirmative consent).

45. H.R. 9619, § 2(a)(4).

46. Press Release, Elec. Privacy Info. Ctr., EPIC Applauds Introduction of Privacy-Centered Federal Chatbot Bill (July 9, 2026), <https://epic.org/press-release-epic-applauds-introduction-of-privacy-centered-federal-chatbot-bill> [hereinafter EPIC Press Release] (praising the bill as "privacy-centered" legislation that protects users from exploitative data practices).

govern, in documented instances targeting people who had criticized the administration or documented immigration enforcement.⁴⁷ Google, Meta, and Reddit complied with at least some of these demands; in one case Google turned over a student’s identifying and financial information to Immigration and Customs Enforcement on a subpoena no judge had approved, and without the advance notice the company’s own policy promised.⁴⁸ Those demands sought subscriber and identifying data rather than the contents of conversations, and the bill’s clause reaches the conversation itself, a category that carries a higher bar even under existing law. The lesson is about provider behavior all the same. A limit that admits government access only on a wiretap warrant is strong on its face, and only as strong as the provider’s willingness to insist on it, and the conduct of these providers under this administration has been to surrender user data on far less. The bill’s answer is a suit the user must discover and bring after the disclosure is already made, often against a gag order, a thin guard for a protection that depends on the provider refusing in the first place.

B. What the Safety Section Builds

The same bill then builds, in a different section, an apparatus that pulls against its privacy sections. The introduced text bars a provider from retaining a user’s chat log “for longer than 5 years, unless retention is necessary to comply with this Act or otherwise required by law.”⁴⁹ It then requires, “[n]ot less frequently than monthly,” that a provider “assess the artificial intelligence chatbot for the risk of any covered harm, emotional dependence, or compulsive usage, according to metrics set forth in rules promulgated by the Commission.”⁵⁰ The assessment

47. *Homeland Security Is Trying to Force Tech Companies to Hand Over Data About Trump Critics*, TechCrunch (Feb. 3, 2026), <https://techcrunch.com/2026/02/03/homeland-security-is-trying-to-force-tech-companies-to-hand-over-data-about-trump-critics/>; *Open Letter to Tech Companies: Protect Your Users From Lawless DHS Subpoenas*, Elec. Frontier Found. (Feb. 2026), <https://www.eff.org/deeplinks/2026/02/open-letter-tech-companies-protect-your-users-lawless-dhs-s> (observing that an administrative subpoena compels nothing absent a court order).

48. *Google Sent Personal and Financial Information of Student Journalist to ICE*, TechCrunch (Feb. 10, 2026), <https://techcrunch.com/2026/02/10/google-sent-personal-and-financial-information-of-student-journalist->

49. H.R. 9619, § 2(a)(1)(F). The model bill on which the introduced text is based set this ceiling at ten years; the introduced text halves it. *See* Model Bill, *supra* note 3.

50. *Id.* § 2(c)(1); *see id.* § 5(7) (defining “Commission” to mean the Federal Trade Commission).

is framed at the level of the chatbot rather than the individual conversation, so the text does not on its face order a scan of each user's stored logs. But its structure still pulls toward retention. Because the metrics do not yet exist and will be written, and rewritten, by the Commission after passage, no provider can know today what a future rule will require it to have measured, or across what history, and the "necessary to comply with this Act" exception to the retention bar invites a provider to hold logs against that uncertainty. A provider that today offers to keep nothing acquires, under this structure, a compliance-flavored reason to begin retaining; and a data-hungry provider already inclined to keep everything can cite the same uncertainty to hold logs toward the five-year ceiling in the name of preemptive compliance. Reinforced by the liability the statute creates, the rational response is to retain rather than delete. Five years of retained conversation, held against a recurring monthly assessment whose target the Commission will define, is not the privacy regime the front of the bill advertises. It tilts toward an inspection regime pointed at the most intimate class of records a person generates.

The cadence also defeats the safety rationale offered for it. A monthly retrospective assessment cannot reach a person in crisis; it arrives weeks after the moment in which intervention could matter. Congress knows how to mandate intervention at the moment of risk when it wants to; a companion measure in the same Congress requires providers to monitor minors' companion chatbot interactions for suicidal ideation and to notify a parental account when such an interaction occurs.⁵¹ Whatever else a monthly sweep of retained logs accomplishes, rescue is not it. The mechanism fits surveillance and does not fit safety.

C. The Deferred Target

What the monthly assessment measures is defined only in part. The bill names the harm categories, covered harm, emotional dependence, and compulsive usage, and defines each.⁵²

51. *Cf.* Children Harmed by AI Technology Act, S. 2714, 119th Cong. (2025) (requiring covered entities to monitor companion chatbot interactions of minors for suicidal ideation and to inform the parental account of any such interaction). Part IX.C addresses why a live monitoring mandate raises its own grave problems; the point here is narrower: a monthly cycle is not a crisis response on any account.

52. H.R. 9619, §§ 5(8), (9), (12).

But it defers the metrics, the operative content of the assessment, to rules the Commission will promulgate describing “the metrics that each artificial intelligence chatbot provider is required to use to assess and publish any risk” of those harms.⁵³ The statute thus permits the retention, requires the recurring assessment, and hands the metrics that give it content to the Commission, after passage. Nothing in the text confines those metrics. One Commission could write them narrowly; another could write them to reach disfavored speech; the statute would tolerate both, because it fixed the categories and left the measure open.

The identity of the rulemaker makes the delegation concrete rather than hypothetical. The office that will write the metrics is the Federal Trade Commission,⁵⁴ and the current Commission has already shown how it approaches delegated authority over AI outputs. On July 7, 2026, it proposed a policy statement declaring that an AI provider which steers its system’s outputs toward undisclosed objectives, including to comply with a state anti-discrimination law, may commit a deceptive act under Section 5 of the FTC Act, and asserting that such a state law is impliedly preempted where it conflicts with that federal scheme.⁵⁵ A statute that leaves the definition of scannable harm to that Commission, or to any Commission of any administration, has not written a safety standard. It has written a blank instrument and handed the pen to an office already using its AI authority to press against civil-rights protection.

D. The Sign on the Door

The transparency section completes the inversion. Among the representations a provider is forbidden to make is any “represent[ation] that the input data or chat log of a user is

53. *Id.* §§ 2(c)(1), 3(a)(3).

54. H.R. 9619, §§ 2(c)(1), 3, 5(7).

55. Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, 91 Fed. Reg. 41,638 (proposed July 7, 2026) (Matter No. P264200; Docket No. FTC-2026-0859), <https://www.federalregister.gov/documents/2026/07/07/2026-13628/policy-statement-concerning-the-suppression-of-accuracy-in-artificial-intelligence-systems> (grounding the theory in Section 5 and naming Colorado’s Artificial Intelligence Act as an example of a state law that may pressure providers to alter outputs); *see* Exec. Order No. 14,365, 90 Fed. Reg. 58,499 (Dec. 11, 2025) (directing the Commission to address how state laws requiring alterations to AI outputs conflict with federal law).

confidential.”⁵⁶ The provision is honest in the narrowest sense: under the retention and assessment regime the bill itself constructs, the conversation is not confidential, and saying so would be false. The consequence is not hypothetical. In February 2026, in a ruling it called a question of first impression, the Southern District of New York held that a criminal defendant’s exchanges with the consumer version of a generative AI tool were neither privileged nor protected work product, in part because the provider’s own terms disclaimed confidentiality and reserved the right to disclose user data to governmental authorities, so the user could have had no reasonable expectation that the exchange was private.⁵⁷ The court left open whether an enterprise tool carrying genuine confidentiality terms might be treated differently.⁵⁸ This bill closes that door. By forbidding every provider from representing that a chat is confidential, it removes the single feature on which a privilege-preserving tool could be built, and it forecloses that path for every user, not only those already under investigation. But observe what the statute has done. The one place a person might bring a thought too dangerous or too shameful to say to another human being is, by operation of this bill, a place with a mandated sign on the door announcing that nothing said inside is private. The bill does not resolve the tension between confidentiality and protection. It resolves it against confidentiality, in every case, for every user, and requires the user to be told so.

Taken together, the privacy sections and the safety section are not complementary halves of one design. The privacy sections bolt the front door against advertisers, data brokers, and warrantless government demands. The safety section props open a back door: five years of retention, a recurring assessment, and metrics to be named later by the Commission. A reader who stops at the privacy sections sees a privacy bill; the endorsements the bill has drawn from the

56. H.R. 9619, § 2(b)(1)(A).

57. *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026) (Rakoff, J.) (holding that documents a defendant generated using a consumer generative AI tool, in connection with a pending criminal prosecution and without counsel’s direction, were protected by neither the attorney-client privilege nor the work product doctrine, and reasoning that the provider’s privacy policy, which reserved the right to disclose inputs to third parties including governmental regulatory authorities, defeated any reasonable expectation of confidentiality).

58. *Id.* (describing the question as one of first impression nationwide, and declining to reach whether an enterprise-grade tool with binding confidentiality terms would change the analysis).

privacy community suggest many readers stopped there.⁵⁹ The architecture rewards that partial reading. That is what makes it dangerous.

E. The Consent That Cannot Be Given

Everything the privacy title does for the adult user runs through a single gate. A provider may not process an adult's chat log for training "unless the artificial intelligence chatbot provider first obtains affirmative consent."⁶⁰ The Act then defines affirmative consent, and the definition is where the provision comes apart.

Affirmative consent means a clear affirmative act signifying the "freely given, specific, informed, and unambiguous authorization of a user."⁶¹ "Informed" in that sequence is not a courtesy. It is a condition of the definition. Where a person cannot be informed, no act that person performs is affirmative consent within the meaning of this Act, and there is nothing for a provider to obtain. The Act then elects the vehicle: the request must include "a written description, in easy-to-understand language, of the act or practice for which the consent of the user is sought,"⁶² and must be "made in a manner reasonably accessible to and usable by a user with a disability (as defined in section 3 of the Americans with Disabilities Act of 1990 (42 U.S.C. 12102))."⁶³

Follow the citation the drafters chose. Section 12102 defines disability as an impairment that substantially limits a major life activity, and its enumeration of major life activities names learning, reading, concentrating, thinking, and communicating, to be construed in favor of broad coverage to the maximum extent the statute permits.⁶⁴ The definition the Act imports to describe the class it means to protect is, in substantial part, a definition of substantial limitation in reading and thinking. The vehicle the Act elects to reach that class is a written description. The

59. See EPIC Press Release, *supra* note 13.

60. H.R. 9619, § 2(a)(1)(C)(iii).

61. *Id.* § 5(2)(A).

62. *Id.* § 5(2)(A)(ii).

63. *Id.* § 5(2)(A)(iii). The subparagraph's parenthetical is unclosed in the introduced text.

64. 42 U.S.C. § 12102(1)(A), (2)(A), (4)(A) (2026).

definition and the mechanism sit in consecutive clauses of one paragraph, and the first forecloses the second.

The drafters were looking directly at that list when they wrote it. Section 5(8) defines compulsive usage by its impact on “one or more major life activities of a user, including socializing, sleeping, eating, learning, reading, concentrating, communicating, or working,”⁶⁵ which is the Americans with Disabilities Act’s phrase and substantially its enumeration, imported into a different definition in the same section of the same bill. The words “reading” and “concentrating” appear twice in Section 5, once as the capacities whose impairment defines the protected class and once as the capacities whose disruption defines the statutory harm. Between those two appearances the Act prescribes reading and concentrating as the method by which the protected class is to authorize the use of its data.

1. What the Record Says About Easy-to-Understand Language

The objection cannot be answered by drafting the notice better, because the evidence on drafting the notice better is already in, and it does not support the subsection.

The premise of Section 5(2)(A)(ii) is that linguistic simplification delivers comprehension. Tested directly, it does not, and the failure concentrates in exactly the population the following subparagraph names. In a controlled study, 102 adults with intellectual disabilities were presented with the consent information for an actual research project under three conditions: the information as written, the information broken into sections with comprehension questions after each, and the information accompanied by colour photographs. There was no significant difference in ability-to-consent scores across the three conditions, and of the 85 participants who completed the protocol, five, or 5.9 percent, met the threshold on every element of consent. The authors concluded that providing additional information or reducing memory load are not in themselves sufficient.⁶⁶

65. H.R. 9619, § 5(8).

66. L. Dye, D. J. Hare & S. Hendy, *Capacity of People with Intellectual Disabilities to Consent to Take Part in a Research Study*, 20 J. Applied Rsch. Intell. Disabilities 168, 168–74 (2007), <https://doi.org/10.1111/j.1468-3148.2006.00310.x> (last visited July 14, 2026) [hereinafter Dye et al.].

Set that experimental design beside Section 5(2)(A). The Act requires a clear and conspicuous standalone disclosure, which is sectioning. It requires easy-to-understand language, which is linguistic simplification. It requires accessibility to users with cognitive disabilities, whose established graphical technique is symbol and photographic support. Those are the three conditions Dye and colleagues ran. The bill's consent architecture is the experimental arm of a study that reported a null result twenty years ago.

The null result replicates and generalizes. A meta-narrative review of 42 studies of easy-read and accessible written information for adults with intellectual disabilities found the measured impact disappointing: where comprehension was assessed and adapted text was compared against unadapted text, only one study found an advantage for the adapted version, and the participants who did better were the cognitively more able and the stronger readers, while for those who struggled more, adapting the text did not tend to help.⁶⁷ That is the finding a legislature needs to sit with. The accommodation tracks ability, and the class the statute exists to protect is defined by the absence of the ability the accommodation tracks. The measure helps in inverse proportion to need.

A later study puts a number on the residue and adds something worse. Students with intellectual disability read or watched blog content in original and linguistically simplified versions; mean comprehension was 55 percent regardless of condition, and the participants predicted they would answer 80 percent correctly. Simplification did not help the weaker comprehenders and was counterproductive for the stronger ones.⁶⁸ The gap between 80 and 55 is the finding that matters for a consent statute. A reader who believes he understood and did not will proceed. He is not hesitant, not confused, and not identifiable to the provider as anything other than a satisfied user. Every doctrine downstream that treats the click as evidence

67. Deborah Chinn & Claire Homeyard, *Easy Read and Accessible Information for People with Intellectual Disabilities: Is It Worth It? A Meta-Narrative Literature Review*, 20 Health Expectations 1189, 1189–1200 (2017), <https://doi.org/10.1111/hex.12520> (last visited July 14, 2026) [hereinafter Chinn & Homeyard].

68. Inmaculada Fajardo et al., *Video-Blogs and Linguistic Simplification for Students with Intellectual Disability*, 35 J. Applied Rsch. Intell. Disabilities 1217, 1217–30 (2022), <https://doi.org/10.1111/jar.13016> (last visited July 14, 2026) [hereinafter Fajardo et al.].

of understanding is reading a signal the data show is uncoupled from the state it is supposed to report.

The pattern is not confined to intellectual disability, and it long predates this bill. Patients presented with a simplified oncology consent form comprehended 58 percent of the essential information, against 56 percent for the standard form written at a college-senior reading level, a difference of two points, while preferring the simplified form by nearly two to one and reporting that it was easier to read and less frightening.⁶⁹ The simplification changed how the instrument felt and did not change what the reader knew. That is the whole of what Section 5(2)(A)(ii) can be expected to deliver: an instrument that is more pleasant to sign.

2. Position, Not Content

For the largest share of the affected population the barrier is not what the consent request says. It is where the consent request stands.

The gate is placed between the user and the thing the user opened the application to do, and that placement is not incidental to it; the placement is its function. On one of the leading accounts of attention-deficit/hyperactivity disorder, the condition is organized around delay aversion, a motivational state in which delay itself acquires negative valence and behavior is oriented to escaping it, and the account predicts that where the delay-averse person cannot avoid actual delay he will allocate attention to whatever speeds the perceived passage of time and so escapes the subjective experience of delay.⁷⁰ The choice pattern is specific to the condition: children with ADHD showed a stronger preference for small immediate over large delayed

69. Terry C. Davis et al., *Health Literacy and Cancer Communication*, 52 CA: Cancer J. for Clinicians 134, 134–49 (2002), <https://doi.org/10.3322/canjclin.52.3.134> (last visited July 14, 2026) (reporting also that patients reading at or below an eighth-grade level comprehended less than half of the essential information in either form).

70. Edmund J. S. Sonuga-Barke et al., *AD/HD and the Capture of Attention by Briefly Exposed Delay-Related Cues: Evidence from a Conditioning Paradigm*, 45 J. Child Psych. & Psychiatry 274, 274–83 (2004), <https://doi.org/10.1111/j.1469-7610.2004.00219.x> (last visited July 14, 2026) [hereinafter Sonuga-Barke et al.] (setting out the delay-aversion account and reporting that children with AD/HD, and not controls, showed an attentional bias toward delay-related cues on a dot-probe conditioning paradigm, which the authors describe as the first evidence of a qualitative rather than quantitative difference in attentional style).

rewards than both control children and children with high-functioning autism, and the preference normalized when stimulation was added to the delay interval.⁷¹

Neuroimaging locates the mechanism and, for present purposes, scales it. Adolescents with ADHD shown cues predicting certain delay, as against cues predicting no delay, displayed significantly greater bilateral activation in the amygdala, anterior insula, temporal pole, and dorsolateral and ventromedial prefrontal cortex; amygdala and dorsolateral prefrontal activation correlated with self-reported delay aversion in a delay-dose-dependent manner and statistically mediated the relationship between ADHD status and delay aversion.⁷² The authors observe that the attentional bias to delay cues in ADHD is equivalent to the bias to social threat cues seen in anxious individuals.

That inverts the assumption the bill is built on. A long instrument standing between a person and the thing he came for is not a boring document to a delay-averse reader. It is a delay cue, processed in the structures that process aversive events, and dose-dependently, so that the longer the instrument the stronger the response. The behavioral answer to an aversive delay cue is escape, and the click is the escape. Nothing about the reading level of a document governs a reader who never opens it, because reading level is a property of a text and the barrier here operates before the text is reached.

Which makes Section 5(2)(A)(ii) a remedy pointed the wrong way. Simplification for comprehension proceeds by replacing abstract constructions with concrete ones and adding examples, and it lengthens the instrument. Length raises the delay dose. The subsection's mechanism for the reading barrier deepens the attention barrier, and the Act repairs neither.

71. Inge Antrop et al., *ADHD and Delay Aversion: The Influence of Non-Temporal Stimulation on Choice for Delayed Rewards*, 47 J. Child Psych. & Psychiatry 1152, 1152–58 (2006), <https://doi.org/10.1111/j.1469-7610.2006.01619.x> (last visited July 14, 2026).

72. Jeroen Van Dessel et al., *Delay Aversion in Attention Deficit/Hyperactivity Disorder Is Mediated by Amygdala and Prefrontal Cortex Hyper-Activation*, 59 J. Child Psych. & Psychiatry 888, 888–99 (2018), <https://doi.org/10.1111/jcpp.12868> (last visited July 14, 2026) [hereinafter Van Dessel et al.] (noting also that the attentional bias to delay cues in ADHD is equivalent to the attentional bias to social threat cues observed in anxious individuals).

None of this requires an impaired reader to demonstrate. When researchers instrumented what consumers do with standard-form contracts in practice, the share who accessed the terms at all was vanishingly small, and those who did spent seconds.⁷³ When researchers built a fictitious social network whose terms of service assigned the user's first-born child, the overwhelming majority of participants accepted without registering the clause.⁷⁴ The gate defeats the median reader. Contract law has absorbed that as consumer indifference. It is better read as the position effect operating at population scale in people whose attention is unimpaired, which means the statute's protected class is the population in whom a mechanism already sufficient to defeat everyone else operates with clinical force.

3. The Inference the Instrument Requires

A third failure sits outside both reading and attention, and no accessible-format requirement touches it.

What a consent request discloses is not carried by its sentences. The sentences state a practice. The disclosure is what that practice will license the provider to do later and what the reader will thereby have given up, which is an inference from text to institutional consequence rather than an operation performed on text. Section 5(2)(A)(ii) requires the description of "the act or practice" to be easy to understand. It does not, and a written description cannot, make the consequence explicit, because the consequence is the part the reader is expected to supply.

Controlled work bears on precisely that step. On an implicature comprehension task administered to 66 autistic and 118 non-autistic adults, and controlling for core language ability, autistic participants were 2.56 times more likely to endorse a non-normative reading of an implied meaning and 6.19 times more likely to answer "do not know" when asked whether an

73. Yannis Bakos, Florencia Marotta-Wurgler & David R. Trossen, *Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts*, 43 J. Legal Stud. 1, 1–35 (2014), <https://doi.org/10.1086/674424> (last visited July 14, 2026).

74. Jonathan A. Obar & Anne Oeldorf-Hirsch, *The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services*, 23 Info. Commc'n & Soc'y 128, 128–47 (2020), <https://doi.org/10.1080/1369118X.2018.1486870> (last visited July 14, 2026).

implied meaning was present.⁷⁵ Core language scores did not differ between the groups. The authors read the “do not know” pattern as a preference for certainty and explicit communication rather than an incapacity, since the normative inference was usually made when the response format did not permit the option, and they recommend that non-autistic communicators accommodate the preference for explicit communication rather than treat it as a deficit.⁷⁶

Carry that finding to a clickwrap. The reader decodes every word correctly. The meaning the drafter intended is the one the drafter left implied. Offered the chance to decline the inference, this reader declines it at six times the rate of his neighbor. And the gate offers no such option. It offers Accept.

4. The Answer They Drafted and Spent on Children

The Act knows that a person can be presented with a consent request and be unable to give consent, because it says so one clause above. Section 2(a)(1)(C)(i) supplies a surrogate, barring processing of a minor’s data “without the affirmative consent of the parent or legal guardian of that user.”⁷⁷ Section 2(a)(1)(C)(ii) does not bother with consent at all, prohibiting training on a minor’s chat log outright, with no gate, no surrogate, and no election available to anyone.⁷⁸ Then Section 2(a)(1)(C)(iii) reaches the adult and offers one route and no surrogate.⁷⁹

Both answers to the problem are drafted. Both are in the same subparagraph. Both are spent on children. The adult who cannot be informed is given the one mechanism the Act’s own imported definition forecloses.

The argument is not that the guardian clause should be extended past a birthday. It should not. Disability law has spent two decades reducing plenary guardianship in favor of supported

75. Alexander C. Wilson & Dorothy V. M. Bishop, “*Second Guessing Yourself All the Time About What They Really Mean...* ”: *Cognitive Differences Between Autistic and Non-Autistic Adults in Understanding Implied Meaning*, 14 *Autism Rsch.* 93, 93–101 (2021), <https://doi.org/10.1002/aur.2345> (last visited July 14, 2026) [hereinafter Wilson & Bishop].

76. *Id.* (reporting that a “do not know” response on the original format was associated with roughly 90 percent accuracy when the same item was re-presented without that option).

77. H.R. 9619, § 2(a)(1)(C)(i).

78. *Id.* § 2(a)(1)(C)(ii).

79. *Id.* § 2(a)(1)(C)(iii).

decision-making, in which a person of the individual's own choosing, who knows him, helps meaning arrive and the decision remains his. That reform is correct, and extending Section 2(a)(1)(C)(i) to adults would run against it while still missing every adult who has no surrogate at all. The point is narrower and harder to answer than a drafting suggestion. The drafters held both answers in hand, on one page, and chose a birthday as the line at which the answers stop.

Nor can supported decision-making be conjured at the interface. The supporter is defined by the knowing; a stranger placed in the role is not performing the role. There is no supporter present at a clickwrap, by construction, and a provider that supplied one would be installing the counterparty as the authority on the user's comprehension of the counterparty's own terms. Were the provider to fill that role with a conversational system, the system would itself be a chatbot under Section 5(3), triggering the notice duty of Section 2(b)(2), which must in turn be accessible to users with disabilities, which is the requirement that produced the problem.

5. Two Standards and No Limit

One further defect belongs on the record because it will decide the litigation. The Act states its accessibility duty twice, at two different strengths. The transparency notice "shall be accessible to users with disabilities,"⁸⁰ unqualified. The consent request need only be made in a manner "reasonably accessible to and usable by" a user with a disability.⁸¹ Neither is accompanied by a feasibility limit, and no feasibility limit appears anywhere in the Act.

That is an anomaly rather than a nicety. Every accessibility obligation in American law carries one. Title III requires removal of architectural barriers only where removal is "readily achievable," a term the statute defines as easily accomplishable without much difficulty or expense, and excuses auxiliary aids that would work an undue burden or fundamentally alter the goods or services offered.⁸² Section 504's implementing regulations are built the same way. An unqualified command to make a notice accessible to users with disabilities, with the class defined

80. *Id.* § 2(b)(2)(B).

81. *Id.* § 5(2)(A)(iii).

82. 42 U.S.C. §§ 12182(b)(2)(A)(iii)–(iv), 12181(9) (2026); 28 C.F.R. §§ 36.303, 36.304 (2026).

by reference to Section 12102, is either an obligation no provider can discharge, or it is an obligation a court will narrow by reading in a limit Congress declined to write, and in the interval the Commission will draw the line by rulemaking. That is the deferral this response traced above at Part III.C, arriving now on the accessibility duty, and the two formulations guarantee litigation over whether the difference between “accessible” and “reasonably accessible” was deliberate.

6. What the Common Law Already Does, and What the Act Does Beyond It

The argument to this point is that the Act’s consent gate does not work. The doctrinal consequence is sharper than that, and reaching it requires no new law.

Consent is not a formality that a defective process leaves merely vulnerable. Mutual assent is a condition of formation, and where assent is absent there is no contract to enforce, to defend, or to police.⁸³ The distinction between whether an agreement was formed and whether a formed agreement is valid is settled and load-bearing. The Supreme Court has held that where the dispute concerns contract formation, the dispute is generally for courts to decide, and has expressly distinguished the question whether an agreement was “ever concluded” from the question whether a concluded agreement was illegal when made.⁸⁴

The First Circuit has already carried that distinction to inaccessible terms. In *National Federation of the Blind v. Container Store*, blind customers enrolled in a loyalty program at in-store touch screens they could not operate, and the program’s terms included an arbitration clause. Applying Texas law, the court held that there “can be no mutual assent or meeting of the minds . . . if the one to whom the offer is supposedly made is unaware of the contract’s terms and conditions,” and that the retailer “failed to meet its burden of establishing that an agreement to

83. RESTATEMENT (SECOND) OF CONTRACTS § 17(1) (Am. L. Inst. 1981) (requiring a manifestation of mutual assent to the exchange).

84. *Granite Rock Co. v. Int’l Bhd. of Teamsters*, 561 U.S. 287, 296–97 (2010) (holding that “where the dispute at issue concerns contract formation, the dispute is generally for courts to decide,” and distinguishing “the generally nonarbitral question whether an arbitration agreement was ‘ever concluded’” from the question whether a contract containing an arbitration clause was illegal when formed); *Buckeye Check Cashing, Inc. v. Cardegna*, 546 U.S. 440, 444 n.1 (2006) (reserving the question whether an agreement was ever concluded). Both cases arise under the arbitration statutes, and the formation-versus-validity distinction is drawn there; the distinction itself is a proposition of general contract law that the arbitration cases apply rather than invent.

arbitrate was ever consummated” with those customers.⁸⁵ It then declined to decide whether the clause was unconscionable, because there was no contract to examine for unconscionability.⁸⁶ Formation came first, and formation was the end of it.

The scope of that holding is the honest part, and it is the part that governs here. The court agreed, in terms, that “inability to read is not a defense to contract formation.”⁸⁷ What defeated formation was not that the plaintiffs were blind. It was that they did not know and had no reason to know that any terms applied at all. The duty to read survives *Container Store* whole. A person handed terms he could have read and did not read is bound by them, and the law has said so for a very long time.

That is what makes Section 5(2)(A) remarkable, and the Act did it to itself.

At common law the duty to read is the backstop that holds the whole clickwrap apparatus up. The law imputes to the signer knowledge of terms he had a fair opportunity to acquire, whether or not he acquired them, and the imputation is what converts a click into assent. H.R. 9619 removes that backstop for its own gate. The Act does not require that the description be presented. It requires that the authorization be “informed,”⁸⁸ and it makes informing an element of the definition of the thing the provider must obtain rather than a description of the process by which the provider must try. A provider cannot answer a challenge to a Section 5(2)(A) authorization by saying the description was there to be read, because the Act never asked whether the description was there to be read. It asked whether the authorization was informed. Congress legislated past the duty to read for the transaction it governs, and then supplied no mechanism, no surrogate, and no fallback for the users in whom the condition it wrote cannot be met.

The consequence follows without novelty, and it is not the consequence the drafters would expect. Where the Act’s condition goes unmet, the resulting authorization is not a

85. *Nat’l Fed’n of the Blind v. Container Store, Inc.*, 904 F.3d 70, 84, 86 (1st Cir. 2018) (applying Texas contract law) [hereinafter *Container Store*].

86. *Id.* (resolving the appeal on formation and not reaching the unconscionability challenge).

87. *Container Store*, *supra* note 85, at 84 (agreeing that inability to read is not a defense, and holding that nonetheless “a party cannot enter into a contract to arbitrate when it does not know or have reason to know the basic terms of the offer”).

88. H.R. 9619, § 5(2)(A).

defective consent that a court might police for unconscionability,⁸⁹ nor an act voidable for incapacity on the transaction-specific showing that doctrine demands.⁹⁰ It is not affirmative consent within the meaning of Section 5(2)(A) at all. Which means the provider that trained on the chat log trained without the consent Section 2(a)(1)(C)(iii) requires, and violated the Act.

The provider's file will say otherwise. It will say so in every case, because the record is generated by the same click the evidence above shows to be uncoupled from the state it purports to report,⁹¹ and because the class in which the condition fails is the class least equipped to say that it failed. That is the shape of the failure this response means to name. A protection routed through a condition part of the protected class cannot satisfy does not fail loudly. It reports success, on paper, in the provider's favor, forever.

7. *What Follows*

The conclusion is not that the Act should drop its accessibility language, which is more than most bills of its kind attempt. The conclusion is that the Act imported the Americans with Disabilities Act's definition of disability, elected written comprehension as its consent vehicle, prescribed a simplification remedy the empirical record does not support, ignored the position of the gate entirely, and then withheld from adults both of the mechanisms it drafted for children. For some share of the class it names, the training gate at Section 2(a)(1)(C)(iii) is not a protection. It is a formality that will be performed, recorded, and relied upon, and the record it generates will not mean what the statute says it means. A provider will hold a signed authorization that satisfies every element of Section 5(2)(A) on its face and satisfies the word "informed" in none of them, and the Act supplies no one with standing to say so and no mechanism to find out.

89. RESTATEMENT (SECOND) OF CONTRACTS § 208 (Am. L. Inst. 1981).

90. RESTATEMENT (SECOND) OF CONTRACTS § 15(1) (Am. L. Inst. 1981) (voidability where a person is unable to understand in a reasonable manner the nature and consequences of *the transaction*); *see Ortelere v. Teachers' Ret. Bd.*, 250 N.E.2d 460, 464–65 (N.Y. 1969) (treating capacity as specific to the transaction at issue rather than as a general attribute of the person). The transaction-specific frame matters here: capacity doctrine asks what this person could understand about *this* instrument, so a person who converses fluently, recalls detail, and holds sustained interests may still be unable to understand the nature and consequences of a nine-thousand-word instrument, and nothing in Section 15 is measured by what a person can recite about a subject that engages him.

91. *See* Fajardo et al., *supra* note 68 (comprehension of 55 percent against a self-predicted 80 percent, regardless of condition).

The fix is available and the Act contains it. Where a consent gate cannot be made to work for a class, the answer is the one the drafters already reached for at Section 2(a)(1)(C)(ii): do not route the protection through consent. A flat prohibition on training over chat logs, subject to a genuine opt-in for users who affirmatively seek one out rather than encounter it, protects the reader who cannot be informed without requiring anyone to certify a comprehension that did not occur. Part IX develops the point.

IV. Two Deferred Terms: The Assessment Metrics and Injury

A. The Metrics, Defined Later

The first deferral is the one Part III traced. The bill defines its harm categories, but the metrics of assessment, the operative content of the safety mandate, are delegated to rules the Commission will write after enactment.⁹² The same rulemaking supplies the form of the required disclosures and the categories of information providers must publish. A substantial part of what the bill does will not exist until the Commission writes it. Legislation whose substance arrives later cannot be evaluated now, and the deferral concentrates that power in a single agency.

B. Injury, Left Unspecified

The second deferral is quieter and runs through the liability title. The introduced text imposes liability on a provider for “an injury in fact caused through the use” of its chatbot, even where the provider “exercised all reasonable care in the design and distribution.”⁹³ Its private right of action separately reaches “a person injured by an act or practice in violation” of the Act.⁹⁴ The phrase

92. *See supra* Part III.C; H.R. 9619, § 3 (rulemaking directives covering disclosure form and content, the assessment metrics, the categories of information providers must publish, and rules prohibiting harmful design features for minors).

93. H.R. 9619, § 4(d).

94. *Id.* § 4(c)(1). Here the introduced text departs from the model bill, which had deemed a privacy or transparency violation to “constitute[] an injury in fact” regardless of concrete harm; the introduced text drops that deeming clause and requires an injury from the violation, a change that removes one standing vulnerability but leaves the scope of “injury” open.

“injury in fact” borrows the standing vocabulary of Article III, but the Act nowhere says what injuries count for its own damages, and its definitions section does not supply the term.⁹⁵

The omission is not housekeeping. Injury is the load-bearing word of the liability regime. Read narrowly, as physical harm, the regime is modest. Read broadly, to include emotional distress, reputational harm, wasted time, or a resolved episode of distorted belief, the regime is far larger, and Part VI.B shows that the broad reading would invert the remedial architecture of federal disability law. A no-fault liability provision keyed to an “injury in fact” the statute never scopes has no stated floor and no stated ceiling for compensable harm; its reach will be discovered in litigation rather than legislated in text. Between the assessment metrics the Commission will define and the injury the statute leaves open, the two terms on which the bill’s coercive force most depends are both unwritten.

V. A Liability Regime Built for a Different Machine

The liability provision imposes, as Part IV noted, liability for any injury in fact without regard to the care the provider exercised.⁹⁶ It does not use the word “product,” but its structure is the structure of strict products liability: recovery for a harm the system caused even where the maker took all reasonable care, on the model of section 402A and its successors.⁹⁷ The drafters plainly intend the analogy to the defective charger or the failed airbag.

The analogy fails at the machine. A toaster that catches fire embodies a defect a careful manufacturer could have prevented; strict liability for such products allocates the cost of preventable defects to the enterprise best positioned to prevent them. A generative language model is not that kind of artifact. Its outputs are probabilistic; no level of care guarantees the next response, and the same properties that make the systems useful, open-ended engagement with any prompt, are the properties that produce the occasional harmful output. Empirical

95. *See* H.R. 9619, § 5 (definitions, omitting “injury”).

96. H.R. 9619, § 4(d).

97. RESTATEMENT (SECOND) OF TORTS § 402A (Am. L. Inst. 1965).

analysis of one extensively documented failure found the model agreeing with the user in more than 85% of its messages in a sampled stretch of the conversation, a behavioral disposition rather than a discrete manufacturing flaw.⁹⁸ A rule that holds the builder liable for any injury notwithstanding all reasonable care, applied to a system that cannot guarantee its output, is not a care-inducing standard. It leaves a careful builder permanently exposed, which in practice forces either restrictions severe enough to break the product or withdrawal of the product, a prohibition wearing liability's clothing.

The deeper point is what the provision reveals about the drafting room. The bill was written by privacy and consumer protection lawyers, and its privacy sections show their skill.⁹⁹ What the text nowhere shows is the participation of anyone versed in machine learning. The evidence is in the provisions: a statute that regulates a screen format rather than a model, builds privacy and safety sections that contradict one another without acknowledgment, defers its safety metrics because it cannot specify them, and imposes a liability rule that treats a probabilistic language system like a wired appliance. Those are the fingerprints of a drafting process that never engaged the technology it regulates, and they compound each of the failures this Commentary traces.

VI. Remedies That Do Not Reach the Injured

A. A Remedy Few Can Lift

The bill's enforcement title authorizes FTC and state enforcement and a private right of action. The private action lets an injured person recover actual damages or, for the privacy and assessment violations, up to \$10,000 per violation; up to \$10,000 in total for the transparency violations; and, for a violation of the minors' design-safety duty, statutory damages of at least

98. *Ex-OpenAI Researcher Dissects One of ChatGPT's Delusional Spirals*, TechCrunch (Oct. 2, 2025), <https://techcrunch.com/2025/10/02/ex-openai-researcher-dissects-one-of-chatgpts-delusional-spirals/> [hereinafter Adler Analysis] (reporting that classifiers applied retroactively to a 200-message sample found more than 85% of the model's messages exhibited unwavering agreement with the user).

99. See Tepper, *supra* note 3 (identifying the model bill's authors as counsel at the Consumer Federation of America, EPIC, and Fairplay).

\$50,000, rising to at least \$250,000 where the violation results in a covered harm and at least \$100,000 where it results in emotional dependence or compulsive usage, with the award subject to increase up to fivefold for willful or knowing conduct and with costs and attorney's fees to a prevailing plaintiff.¹⁰⁰ On paper, the fee-shifting and statutory floors are the standard machinery for making small consumer claims viable.

The paper machinery should be measured against the most instructive precedent available: the private right of action that federal law has given, for more than three decades, to the population most exposed to technological exclusion. Title III of the Americans with Disabilities Act confers a private right of action on people with disabilities denied full and equal enjoyment of public accommodations,¹⁰¹ and limits the private remedy to the relief of the Civil Rights Act of 1964's public accommodations title, injunctive relief without damages.¹⁰² That right goes substantially unexercised, because the burden of enforcement falls on the injured individual, who must find counsel, front the cost, and endure the process. A remedy that exists in text and is out of reach in practice is the documented condition of the strongest analogous shield in federal law. The new private right of action inherits the same weight-bearing structure: an injured individual, a lawsuit, and everything a lawsuit demands. Nothing in the bill lightens it.

B. The Emotional Distress Inversion

If the bill's undefined injury is read to include emotional or psychological harm, the statute does something its drafters may not have noticed: it grants chatbot users a remedy that federal disability law categorically denies to the people that law protects.

In *Cummings v. Premier Rehab Keller*, the Supreme Court held that "[e]motional distress damages are not recoverable in a private action to enforce either the Rehabilitation Act of 1973 or the Affordable Care Act."¹⁰³ The statutes at issue are the anti-discrimination guarantees

100. H.R. 9619, § 4(c).

101. 42 U.S.C. § 12182(a).

102. 42 U.S.C. § 12188(a)(1) (adopting the remedies and procedures of 42 U.S.C. § 2000a-3(a) for private Title III actions).

103. *Cummings v. Premier Rehab Keller*, 596 U.S. 212 (2022).

that protect disabled people against entities receiving federal funds.¹⁰⁴ The Court reasoned that Spending Clause legislation operates like a contract, and a remedy is available “only if the funding recipient is on notice that, by accepting federal funding, it exposes itself to liability of that nature.”¹⁰⁵ Because “emotional distress is generally not compensable in contract,” recipients were not on notice, and the damages are unavailable.¹⁰⁶ Justice Breyer, dissenting, would have applied the traditional contract exception under which emotional distress damages lie where serious emotional disturbance was a particularly likely result of the breach, observing that intentional invidious discrimination is exactly such a case.¹⁰⁷

Cummings rests on the absence of notice in statutes that are silent about remedies. It does not govern a statute that expressly authorizes damages, and this bill authorizes damages expressly. The consequence is an inversion. A chatbot user who suffers emotional harm would recover under the bill’s plain terms. A disabled person subjected to intentional discrimination by a federally funded program recovers nothing for the same species of harm, no matter how egregious the discrimination, unless the Supreme Court revisits *Cummings* or Congress overrides it. The bill’s rule of construction preserves both regimes side by side,¹⁰⁸ which produces the final indignity: a disabled person harmed emotionally through a chatbot may find it easier to plead the harm as a consumer claim under this statute than as the civil rights violation it is. A consumer protection bill for chatbot users would thereby out-remedy the civil rights architecture built for the most vulnerable, and the distance between the two measures how far the disability remedy has been hollowed. The objection does not reach the remedy itself. A remedy for emotional harm is worth having, and chatbot users should have one. The objection reaches the priority. A Congress able to write this remedy for a new and visible class has not written it for the disabled,

104. 29 U.S.C. § 794; 42 U.S.C. § 18116.

105. *Cummings*, 596 U.S. at 220 (quoting *Barnes v. Gorman*, 536 U.S. 181, 187 (2002)).

106. *Id.* at 221.

107. *Id.* at 232 (Breyer, J., dissenting) (quoting RESTATEMENT (SECOND) OF CONTRACTS § 353, at 149 (Am. L. Inst. 1979)).

108. H.R. 9619, § 6 (“Nothing in this Act or any regulation promulgated under this Act may be construed to affect any right, cause of action, remedy, presumption, liability, or defense available at law or in equity, including any anti-discrimination, consumer protection, labor, or civil rights law.”).

though it could close the *Cummings* gap whenever it chose, and the bill compounds the disparity by rewarding a plaintiff who recasts a civil rights injury as a consumer one. Whether courts would read the bill's injury to reach emotional harm is unsettled; products liability doctrine has centered physical injury, and the reading should be developed with care. The exposure is stated here because the drafters never defined the word that decides it.¹⁰⁹

C. The Bystander Gap

The bill's remedy runs to users, and only to users.¹¹⁰ The most severe documented harms in which chatbots have figured fell on people who never used one.

On April 17, 2025, a gunman opened fire near the student union of Florida State University, killing Tiru Chabba, a regional vice president at the campus food service vendor Aramark, and Robert Morales, a Florida State dining coordinator, and wounding six others; according to state authorities and the ensuing federal complaint, the shooter had consulted ChatGPT while planning the attack, and the system supplied information about weapon and ammunition selection, the busiest hours at the targeted location, and the observation that media attention increases when children are involved.¹¹¹ Florida's Attorney General has opened a criminal investigation into OpenAI's role.¹¹²

On February 10, 2026, an eighteen year old shooter killed eight people in Tumbler Ridge, British Columbia, among them five students aged twelve and thirteen and an educational assistant, and wounded more than two dozen, in Canada's deadliest school shooting in decades; the complaints allege that OpenAI's own safety personnel had flagged the shooter's account in

109. *See supra* Part IV.B.

110. *See supra* note 8 and accompanying text.

111. *Lawsuit Accuses ChatGPT of Helping Gunman Plan FSU Mass Shooting*, PBS News (May 11, 2026), <https://www.pbs.org/newshour/nation/lawsuit-accuses-chatgpt-of-helping-gunman-plan-fsu-mass-shooting> (reporting the state disclosures and the widow's federal suit); *Lawsuit Says ChatGPT Told FSU Shooter that Targeting Children Would Bring More Attention*, NBC News (May 11, 2026), <https://www.nbcnews.com/news/us-news/openai-sued-chatgpts-alleged-role-guiding-fsu-shooter-rcna344443> (detailing the complaint's allegations, including the timing advice and the peak-hours response).

112. *Florida AG Launches Criminal Investigation into ChatGPT over FSU Shooting*, NPR (Apr. 21, 2026), <https://www.npr.org/2026/04/21/nx-s1-5793967/florida-openai-investigation-mass-shooting-fsu>.

June 2025 as presenting a credible threat and recommended contacting police, that leadership declined, and that the shooter simply opened a new account.¹¹³ Seven families have sued OpenAI and its chief executive personally in federal court in California, with counsel signaling as many as two dozen more suits,¹¹⁴ and the Province of British Columbia has announced that it is preparing its own action.¹¹⁵

Not one of those victims was a user. The two men did not use the chatbot; the murdered children did not use the chatbot; the wounded twelve year old in intensive care did not use the chatbot. In every one of these cases, the user was the attacker. Under the People-First Chatbot Act, the statute's signature remedy gives these plaintiffs nothing, because the harm the chatbot facilitated did not fall on the person at the keyboard.

The Tumbler Ridge families face a second and independent exclusion. As foreign nationals injured abroad, they would stand outside the bill's private remedy even if Congress redrew it to reach bystanders, under the settled presumption that federal statutes do not reach injuries suffered outside the United States absent a clear congressional statement, a statement this bill nowhere attempts.¹¹⁶ Their claims proceed in an American court against an American defendant under ordinary tort law, which carries no comparable barrier. The division of proof between the two attacks is therefore exact. Florida State demonstrates on the statute's own

113. *Families of Canadian Mass Shooting Victims Sue OpenAI, CEO Altman in U.S. Court*, CNBC (Apr. 29, 2026), <https://www.cnbc.com/2026/04/29/families-of-canadian-mass-shooting-victims-sue-openai-ceo-altman.html> (reporting the June 2025 internal flag, the recommendation to contact police, and the decision not to); *Families Sue OpenAI over Tumbler Ridge Mass Shooter's Use of ChatGPT*, NPR (Apr. 29, 2026), <https://www.npr.org/2026/04/29/nx-s1-5798896/tumbler-ridge-mass-shooting-chat-gpt-lawsuit> (describing negligence and defective design theories); *Families of Tumbler Ridge Shooting Victims Sue OpenAI and CEO Sam Altman*, CNN Business (Apr. 29, 2026), <https://www.cnn.com/2026/04/29/tech/openai-tumbler-ridge-canada-shooting-lawsuits> (identifying the deceased and the claims).

114. Zachary Folk, *OpenAI and Sam Altman Could Face Dozens More Lawsuits over School Shooting in British Columbia, Lawyer Says*, Forbes (Apr. 29, 2026), <https://www.forbes.com/sites/zacharyfolk/2026/04/29/openai-and-sam-altman-could-face-dozens-more-lawsuits-over-school-shooting-in-british-columbia-lawyer-> (reporting counts of negligence, product liability, wrongful death, and aiding and abetting).

115. *Canadian Province Preps OpenAI Lawsuit over Alleged ChatGPT-Linked Shooting*, Al Jazeera (July 7, 2026), <https://www.aljazeera.com/news/2026/7/7/canadian-province-sues-openai-over-alleged-chatgpt-linked-shoot>

116. See *Morrison v. Nat'l Austl. Bank Ltd.*, 561 U.S. 247 (2010) (holding that federal statutes are presumed not to apply extraterritorially absent a clear congressional indication); *RJR Nabisco, Inc. v. European Cmty.*, 579 U.S. 325 (2016) (holding that a private right of action under a federal statute requires a domestic injury even where the statute reaches some foreign conduct).

soil that the remedy stops at the keyboard. Tumbler Ridge demonstrates the pattern the statute ignores, and shows that the common law this bill duplicates and narrows is the only American law that was ever going to reach those families at all.

The common law knows better. Strict products liability has extended to bystanders for more than half a century; the California Supreme Court held in *Elmore v. American Motors Corp.* that “the doctrine of strict liability in tort is available in an action for personal injuries by a bystander against the manufacturer and the retailer,” reasoning that bystanders “should be entitled to greater protection than the consumer or user where injury to bystanders from the defect is reasonably foreseeable,” because the bystander has no opportunity to inspect the product or choose its maker.¹¹⁷ The pending suits proceed on exactly that inheritance: wrongful death, negligence, failure to warn, and products liability theories that have never required the victim to be the purchaser or the user. Those theories face genuine obstacles, including the interactive computer service immunity and First Amendment defenses the industry has raised,¹¹⁸ but they exist, they reach bystanders, and they are being litigated now.

D. Redundant with Better Law

The comparison completes itself. The accountability this bill offers is a private right of action for users, layered onto strict liability language a court may not sustain. The accountability the victims of the worst chatbot involved harms are pursuing is the ordinary common law of products and negligence, which reaches bystanders, reaches wrongful death, and is in active use against the industry today.¹¹⁹ The no-fault liability this bill would create runs only to a harmed user, while the plaintiffs who most need it, the bystanders, are outside the statute’s protected class. A measure that duplicates the accountability courts are already providing, narrows it to the keyboard, defers its safety substance to an unwritten rulemaking, and adds a retention and

117. *Elmore v. Am. Motors Corp.*, 451 P.2d 84, 89 (Cal. 1969).

118. 47 U.S.C. § 230; see *Character.ai Bans Open-Ended Chatbot Conversations for Users Under 18*, Rolling Stone (Oct. 29, 2025), <https://www.rollingstone.com/culture/culture-news/character-ai-sued-teens-suicide-banned-minors-chatbots-1235456426/> (noting the defendant’s reliance on Section 230 and free speech protections in pending litigation).

119. See *supra* notes 41–45 and accompanying text.

scanning apparatus in exchange, is not protection. It is the appearance of action on a problem the existing law is already addressing, staged for an audience that will not read past the title.

VII. Marketing a Condition the Statute Does Not Name

The coalition's public case for the bill lists "AI psychosis" among the harms the legislation answers.¹²⁰ The operative text answers no such thing. The introduced bill defines its gravest harm category, "covered harm," as "death, a suicide attempt, or a psychiatric emergency that results in urgent medical treatment,"¹²¹ and the word "psychosis" appears in no provision. The nearest the text comes is a "psychiatric emergency," a clinically distinct and far narrower thing. The gap between the marketing and the text is a fair measure of the enterprise: the term does the recruiting, and the statute does not carry it.

The term itself does damage independent of the gap. AI psychosis is not an established diagnosis; it appears in no diagnostic manual. The clinical category it borrows is defined by fixity: the diagnostic manual characterizes delusions as fixed beliefs that do not yield to conflicting evidence.¹²² The most publicized case offered for the coinage illustrates the difference rather than the condition. A Canadian user with no psychiatric history was drawn by weeks of sycophantic reinforcement into believing he had discovered a revolutionary mathematical theory; he questioned the belief repeatedly throughout, and it collapsed at the first contradiction from a second system, after which he retained no attachment to it and sued the developer.¹²³ A belief that is repeatedly doubted while held and abandoned on first contradiction

120. Fairplay Press Release, *supra* note 4 (asserting that chatbots "can pose a risk of severe harms to young children, including suicide, AI psychosis, compulsive use, and unhealthy emotional dependency").

121. H.R. 9619, § 5(9).

122. AM. PSYCHIATRIC ASS'N, DIAGNOSTIC AND STATISTICAL MANUAL OF MENTAL DISORDERS (5th ed., text rev. 2022) (defining delusions in the chapter on Schizophrenia Spectrum and Other Psychotic Disorders).

123. *They Thought They Were Making Technological Breakthroughs. It Was an AI-Sparked Delusion*, CNN Business (Sept. 5, 2025), <https://www.cnn.com/2025/09/05/tech/ai-sparked-delusion-chatgpt> (reporting the user's repeated expressions of doubt and the belief's collapse upon contradiction by a second chatbot); Anthony Milton, *Man vs. Machine*, Toronto Life (Mar. 3, 2026), <https://torontolife.com/deep-dives/man-vs-machine-chatgpt-delusion-now-hes-suing-openai/> (reporting the absence of psychiatric history, more than 50 requests for reassurance during the episode, and the subsequent lawsuit).

fails the fixity that defines the clinical category. What the documented cases show is real and serious: systems disposed toward agreement can reinforce, amplify, and fail to interrupt distorted or dangerous thinking in the people who bring it to them.¹²⁴ What they do not show is a machine that produces psychosis in the ordinary user, and the difference is not pedantry. People live with psychotic disorders, and for them the word names a condition that structures their lives. Advocacy that stretches the word into a slogan dilutes the condition it borrows, misdescribes the sufferers it recruits, and stigmatizes the population that carries the diagnosis, all in the service of a bill whose text never engages the subject. A measure sold as protection for the vulnerable should not fund its marketing with the vocabulary of the most vulnerable.

The bill's actual answer to the bonding harms its supporters describe is a disclosure: notice that the user is interacting with a chatbot rather than a human, delivered before the first output, every hour thereafter, and whenever the user asks.¹²⁵ The notice has a narrow and real function against disguised bots and false credentials. Against the harm the coalition is loudest about, it is cosmetic. The attachment the advocates describe is not produced by a mistaken belief that the machine is human; it is produced by fluency, memory, and inexhaustible responsiveness, and it survives full knowledge of the machine's nature. Meanwhile, the emotionally central example in the public case, the minor bonded to a companion persona, addresses a door already closing: the platform at the center of those cases removed open-ended chat for users under eighteen effective November 2025.¹²⁶ The live vector, the general purpose assistant embedded in everything, is exactly what the bill's chatbot-only coverage leaves loose.

124. Adler Analysis, *supra* note 28.

125. H.R. 9619, § 2(b)(2) (requiring notice that the user is interacting with a chatbot before any output, every hour thereafter, and each time a user asks whether they are interacting with a real person).

126. Character.AI, *Taking Bold Steps to Keep Teen Users Safe on Character.AI*, Character.AI Blog (Oct. 29, 2025), <https://blog.character.ai/u18-chat-announcement/>; Character.AI, *An Update on Changes to Our Under-18 Experience*, Character.AI Blog (Nov. 21, 2025), <https://blog.character.ai/an-update-on-changes-to-our-under-18-experience/> (describing the staged removal of open-ended chat for users under eighteen beginning November 24, 2025).

VIII. Restriction Without Capability

Every mechanism in the bill restricts, warns, or punishes. The provider is restricted in its data practices, compelled to disclose, assessed monthly, and exposed to liability; the user is warned hourly and handed a cause of action. Nothing in the bill, and nothing in the broader campaign behind it, invests in the capability of the people it claims to protect. The coalition's own program alongside the bill consists of an advisory against AI toys, a demand that a video platform stop recommending AI generated content to children, and a call for a five year pause on generative AI in schools.¹²⁷ Restriction, demand, and pause: the program contains no plank that teaches a parent how these systems fail, no plank that equips a child to recognize a machine's flattery, no plank that makes any user more capable of meeting the technology than they were before.

The omission is a theory of change, and it is the wrong one. A shield-only strategy leaves the protected population permanently dependent on the shield, and Part VI.A documented what happens to shields that depend on individual enforcement: they sit unused.¹²⁸ The documented cases point the same direction. The user whose mathematical revelation collapsed at one outside contradiction did not need a statute; he needed to know that these systems agree by disposition and that a claimed discovery is tested by experts rather than by the machine that proposed it. That knowledge is teachable, it is cheap relative to litigation, and it is the one intervention that reaches the largest class of cases, the person of ordinary judgment fooled by fluent agreement. No provision of this bill attempts it.

IX. A Better Path

The goal, safer systems and real accountability, is achievable with the same legislative energy. The elements follow from the failures, and none is speculative; each has a working model in law already enacted.

127. Fairplay Press Release, *supra* note 4 (describing the “Childhood Made by Humans” campaign).

128. *See supra* Part VI.A.

A. Regulate by Consequence Rather than Format

Extend coverage to automated systems that make or materially influence decisions about a person's employment, education, housing, benefits, or liberty, whether or not the system converses.¹²⁹ The harm lives in the decision and the person it lands on, and a person harmed through a covered system should hold the claim whether or not they were its user.

Working legislatures have already drafted to this unit. Colorado's current statute regulates "automated decision-making technology" used to materially influence a consequential decision concerning education, employment, housing, financial services, insurance, health care, or government services, whatever the interface, and expressly excludes "consumer-facing conversational tools that are not intended for consequential decision-making."¹³⁰ Illinois, effective January 1, 2026, makes it a civil rights violation for an employer to use artificial intelligence, generative systems included, that has the effect of subjecting employees to discrimination on the basis of a protected class in recruitment, hiring, promotion, discipline, discharge, and the other incidents of employment, with the duty attaching to the decision rather than to the tool's format.¹³¹ Colorado watches the decision and disregards the interface; this bill watches the interface and disregards the decision. Under the Illinois statute, the rejected applicant of Part II holds a claim; under this bill, that applicant holds nothing.

129. *See supra* Part II.

130. S.B. 26-189 (Colo. 2026) (signed May 14, 2026; effective Jan. 1, 2027), <https://leg.colorado.gov/bills/sb26-189>; *see Colorado Enacts Artificial Intelligence Replacement Law*, Seyfarth Shaw LLP (May 22, 2026), <https://www.seyfarth.com/news-insights/colorado-enacts-artificial-intelligence-replacement-law.html> (describing the covered technology definition and the conversational-tool exclusion); *Colorado Enacts Revised AI Law*, Norton Rose Fulbright (May 2026), <https://www.nortonrosefulbright.com/en-us/knowledge/publications/18733d31/colorado-enacts-revised-ai-law> (listing the consequential decision domains and the consumer rights to notice, explanation, correction, and meaningful human review).

131. H.B. 3773, 103d Gen. Assemb. (Ill. 2024) (effective Jan. 1, 2026), <https://www.ilga.gov/ftp/legislation/103/BillStatus/HTML/10300HB3773.html> (amending the Illinois Human Rights Act to reach artificial intelligence, defined to include generative systems, that has a discriminatory effect, and barring zip codes as a proxy for protected classes); *see Illinois Anti-Discrimination Law to Address AI Goes into Effect on 1 January 2026*, Nat'l L. Rev. (2025), <https://natlawreview.com/article/illinois-anti-discrimination-law-address-ai-goes-effect-1-january-2026>.

B. Minimize Data Rather than Retain and Scan It

Require prompt deletion of chat logs as the default, and strike the mandate to assess retained conversations against criteria written later by the government.¹³² Privacy protected at the front door and surrendered at the back is not privacy. Minimization is the first principle of modern privacy law, and the European Union names it as one, confining personal data to what is adequate, relevant, and limited to what is necessary for the purpose.¹³³ A safety mandate that requires the opposite forfeits the name.

Two constraints follow from Part III.E, and they bind this proposal exactly as they bind the bill.

The first is that the protection cannot ride on a gate the user must pass to reach the thing he came for. That is the structure Part III.E describes, and it is a structure rather than a word. The defect is not consent. It is interposition: an instrument on the only road, the object on the far side, the click as the toll. A person who cannot be informed meets that gate whether or not he wants what is behind it, because what is behind it is the application.

An election the user seeks is a different object. He navigates to it. Nothing he was already trying to reach sits behind it. If he never finds it, the default has already protected him. The failure modes traced above do not travel to it: nothing is interposed, so there is nothing to escape; the delay cue does not fire on a delay the person went looking for; and the reader who cannot be informed is never asked to authorize anything, because he does not arrive and does not need to. The bill's own architecture cannot draw this line, because Section 5(2)(A) defines affirmative consent by reference to a request the provider issues rather than one the user seeks.¹³⁴

So: prompt deletion as the default, and, for the user who wants his own logs kept, a retention period he sets himself, in a duration he chooses, from a setting he goes and finds, revocable in one step from the same place, never surfaced in the path of another task, and never

132. *See supra* Part III.B–C.

133. Regulation (EU) 2016/679, art. 5(1)(c), 2016 O.J. (L 119) 1 (“data minimisation”).

134. H.R. 9619, § 5(2)(A) (defining affirmative consent by reference to “a specific request from an artificial intelligence chatbot provider”).

a condition of using the product. The person who cannot navigate to it loses nothing, because the default is the protection. The person who wants it is not held to the reading level of the person who does not.

The second is non-imputation, and it is what holds the first in place. No provision should treat an interface action as evidence that the person understood what the action licensed. Not the click, not the scroll, not the seconds elapsed, not the checkbox, and not the recital that the description was displayed. The common law supplies that imputation through the duty to read, and, as Part III.E shows, this bill withdraws the duty to read for its own gate without saying so and without replacing it. A statute in that posture has to finish the thought: say that no inference of comprehension runs from conduct at the interface, and let the protection stand on the default rather than on the record. The practical test is a single question a drafter can ask of any provision. Does the protection survive a plaintiff who proves he understood nothing? If it does not, the statute has not protected the class it named. It has documented them.

C. Keep Safety Detection in the Moment

If detection is mandated, mandate it as an ephemeral intervention in the moment: detect the crisis in the exchange where it is happening, surface help there, and retain nothing, report nothing, and reuse nothing afterward. The disposition that produces these crises is documented in the peer-reviewed literature: assistants tuned on human feedback consistently prefer responses that match the user's beliefs over accurate ones, across every leading system tested, a general and measurable property rather than a mystery.¹³⁵ What is measurable in the laboratory is detectable in the exchange. The capability exists; classifiers applied to the documented reinforcement failure

135. Mrinank Sharma et al., *Towards Understanding Sycophancy in Language Models*, Twelfth Int'l Conf. on Learning Representations (2024), <https://openreview.net/forum?id=tvhaxkMKAn> (demonstrating consistent sycophancy across five leading commercial AI assistants on varied free-form generation tasks and tracing the behavior to the human preference data used in training).

flagged the harmful pattern message by message.¹³⁶ A provider necessarily processes each message to answer it; the legislative choice is what survives the moment. Nothing should.

D. Put the Substance in the Statute

Define the operative terms, injury included, and the core safety expectations in the legislation itself, or set them through a bounded public process that includes machine learning and clinical expertise, rather than delegating the meaning of harm to a single office after passage.¹³⁷

E. Match Liability to the Machine

Replace liability without regard to care with a regime suited to probabilistic systems: a duty of reasonable care in design and deployment, defined safe harbors for providers that implement specified safeguards, and liability for negligent design and for known failure modes left unaddressed.¹³⁸ Hold builders to what care can control, and the standard will induce care rather than withdrawal.

The legislative record already marks the outer boundary. Colorado's original artificial intelligence act imposed a duty of reasonable care on developers and deployers of high-risk systems, cushioned by rebuttable presumptions for providers that maintained risk management programs and impact assessments, a regime far gentler than liability without regard to care.¹³⁹ Even that regime never took effect: it drew a constitutional challenge in which the federal government intervened, a stipulated stay of enforcement in April 2026, and repeal the following

136. Adler Analysis, *supra* note 28 (reporting that open-sourced classifiers, applied retroactively, repeatedly flagged the responses that reinforced the delusion); *cf.* S. 2714, 119th Cong. (2025) (mandating monitoring of minors' interactions with parental notification, a design that converts the confidant into a reporting channel and chills the disclosure it means to catch).

137. *See supra* Part IV.

138. *See supra* Part V.

139. S.B. 24-205 (Colo. 2024), <https://leg.colorado.gov/bills/sb24-205> (requiring developers and deployers of high-risk artificial intelligence systems to use reasonable care to protect consumers from known or reasonably foreseeable risks of algorithmic discrimination, with a rebuttable presumption of reasonable care for those implementing a risk management program and impact assessments).

month in favor of a statute of notice and rights.¹⁴⁰ A Congress proposing liability without regard to care for the same technology is proposing a regime harsher than one that could not survive to its own effective date. Duties of care with defined safe harbors are what the record will bear, and they are enough.

F. Make the Remedy Usable

Lead with public enforcement rather than resting the statute on individual suits, and build accessible complaint and resolution channels that an injured person can invoke without counsel.¹⁴¹ The pattern is familiar in civil rights practice, where an injured person may begin with an administrative charge rather than a federal complaint,¹⁴² and it is the difference between a remedy that exists and a remedy that gets used. Where private actions lie, extend them to any person injured through a covered system, on the bystander principle the common law settled decades ago.¹⁴³

G. Invest in the Capability of the Protected

Capability is already treated as infrastructure by the world's principal AI regulator and by the current administration. The European Union made AI literacy the first obligation of its comprehensive AI regulation to enter application, in force since February 2, 2025, requiring providers and deployers alike to ensure a sufficient level of AI literacy in the people who operate

140. *Colorado's Artificial Intelligence Law Could Be on the Chopping Block*, Littler (updated Apr. 2026), <https://www.littler.com/news-analysis/asap/colorados-artificial-intelligence-law-could-be-chopping-block> (reporting the April 27, 2026 order staying enforcement pending rulemaking); *Colorado Enacts Artificial Intelligence Replacement Law*, Seyfarth Shaw LLP (May 22, 2026), <https://www.seyfarth.com/news-insights/colorado-enacts-artificial-intelligence-replacement-law.html> (describing S.B. 26-189's repeal and replacement of S.B. 24-205 before its effective date).

141. See *supra* Part VI.A; cf. 42 U.S.C. § 12188(b)(2)(B)–(C) (authorizing, in the Attorney General's public action, monetary damages for persons aggrieved and civil penalties of up to \$50,000 for a first violation and \$100,000 for subsequent violations, as periodically adjusted for inflation, relief the private action does not carry).

142. See *Artificial Intelligence in Employment Update: Illinois Requires Notice and Prohibits Discriminatory Impact in Use of AI*, Crowell & Moring LLP (Sept. 2024), <https://www.crowell.com/en/insights/client-alerts/artificial-intelligence-in-employment-update-illinois-requires-notice-and-prohibits-discriminatory-impact-in-use-of-ai> (describing enforcement of the Illinois artificial intelligence provision through the state human rights charge process or a private action).

143. See Elmore, 451 P.2d at 89.

these systems, calibrated to the persons on whom the systems are used.¹⁴⁴ The White House, by executive order, has established a national program of AI literacy resources and teacher training for schoolchildren.¹⁴⁵ The one population no enacted or pending measure equips is the ordinary adult user of a consumer chatbot, the population this bill names as its beneficiary. Fund AI literacy for parents, for children, and for the public, in schools and in communities: how the systems work, where they fail, why their agreement is a disposition rather than a judgment, and how to test what they produce. Accountability for builders and capability for users are halves of one strategy, and a bill that enacts only the first leaves the vulnerable exactly as vulnerable as it found them.

X. Conclusion

The People-First Chatbot Act arrives with the right goal and the wrong instrument. It watches the surface where the technology talks and ignores the machinery where it decides. It bolts the front door of the chat log and builds a back door of retention and scanning, with the target of the scan left to the discretion of whoever holds an office that today has no confirmed occupant. It rests its coercive force on one term the rulemaker will define later and another no one has defined at all. It aims strict liability at a machine that cannot promise its next sentence, hands the injured a remedy the last thirty years of disability enforcement prove will go unlifted, and closes its doors entirely to the bystanders of the two worst attacks in which these systems have figured, plaintiffs the ordinary common law is serving better at this moment.¹⁴⁶ It is marketed on a condition its

144. Regulation (EU) 2024/1689, art. 4 (AI Act); see *AI Literacy – Questions & Answers*, Eur. Comm’n, <https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers> (last visited July 11, 2026) (confirming that Article 4 has applied since February 2, 2025, with national supervision and enforcement from August 2026).

145. Exec. Order No. 14,277, *Advancing Artificial Intelligence Education for American Youth* (Apr. 23, 2025), <https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth/> (establishing a White House task force, directing development of K-12 AI literacy resources, and creating teacher training programs).

146. See *supra* Parts II–VI.

text never names, at the expense of the people who carry the real diagnosis, and it declines to spend one provision making anyone more capable of meeting the technology it fears.¹⁴⁷

Legislation that performed less and protected more is available, and Part IX sketches it. Until something like it is on the table, the accurate name for this bill is the one its structure earns. It is safety theater: the staging of protection for an audience expected to applaud without reading, while nine hundred million confessionals sit retained and scanned, month after month, for whatever the office in power has most recently decided to count as harm.¹⁴⁸

147. *See supra* Parts VII–VIII.

148. OpenAI reported more than 900 million weekly active users of ChatGPT in February 2026, against 400 million a year earlier, 700 million in July 2025, and more than 800 million in October 2025; press accounts in June 2026 reported the application had passed one billion monthly active users. The figure in text is the weekly count and is therefore conservative.