

# Safety Theater

*What the People-First Chatbot Act gets wrong, and what a real chatbot safety law would do*

On July 9, 2026, Representatives Foushee and Casar introduced H.R. 9619 as a comprehensive answer to chatbot harm. The Real Safety AI Foundation opposes the bill as introduced. Read at the provision level, it regulates the conversational interface rather than the technology, bars advertisers from chat logs while permitting those logs to be retained for five years against a monthly assessment whose metrics the FTC will write later, rests its liability on an "injury in fact" the text never defines, and closes its remedy to the bystanders of the gravest chatbot-involved attacks. A response, and the real version.

# The goal is right and overdue. The statute, as drafted, does not meet it.

Introduced July 9, 2026 by Representatives Foushee and Casar, built on the January 2026 model bill from the Consumer Federation of America, EPIC, and Fairplay. This response reads the **introduced House text**, provision by provision.

Every failure traces to one choice: **regulate the conversational surface where harm is most visible, not the technology and consequences where harm is most severe**, and defer the safety substance to a rulemaking that has not occurred.

*The result is legislation that performs protection. The Foundation opposes H.R. 9619 as introduced.*

- 1

**The wrong unit.** It covers the interface; the identical models deciding hiring, housing, benefits, and admissions escape it entirely.
- 2

**The wrong actor.** Every duty runs to the provider who makes a chatbot; the word "deployer" appears nowhere, so the person who aims an agent at a stranger owes nothing while the open-source maintainer answers.
- 3

**Privacy against safety.** It bars advertisers, sales, and warrantless access, then permits five-year retention and a monthly assessment that pushes toward retaining and scanning.
- 4

**Consent that cannot be given.** The training gate demands "informed" consent by written notice, then imports the ADA definition of disability, in substance a definition of limitation in reading and thinking.
- 5

**Terms named later.** The assessment metrics are deferred to the FTC, and "injury in fact," the trigger of liability, is never scoped in the text.
- 6

**A prohibition in liability's clothing.** No-fault liability on systems that cannot guarantee their next output; Colorado repealed a far milder duty before it took effect.
- 7

**No bystander remedy.** The Florida State and Tumbler Ridge victims were not users; their families take nothing from this bill.
- 8

**Marketed on a word it never uses.** Sold on "AI psychosis," a condition the text never names, with nothing invested in the AI literacy of the people it claims to protect.

# What the Bill Gets Wrong

*Seven failures, each traced to the same drafting choice*

- 1 The wrong unit and actor
- 2 The privacy inversion and consent
- 3 Two deferred terms
- 4 The wrong machine
- 5 Remedies that miss
- 6 A condition never named
- 7 Restriction without capability

# It regulates the surface where the machine talks, and ignores the machinery where it decides.

Coverage attaches to systems whose responses "are not fully predetermined" and that accept "open-ended natural-language or multimodal user input." The same models also screen job applicants, rank admissions, score benefits, and inform risk assessments; none converses, so none is covered. **The heaviest obligations land on the most visible deployment; the consequential decision systems go free.**

## IT FAILS INSIDE ITS OWN CATEGORY

A hiring manager asks a covered chatbot whom to reject; the answer discriminates. The "user" is the manager, uninjured. **The rejected applicant never used the chatbot and holds no claim**, and the same structure repeats for the landlord, the caseworker, the teacher.

# The line that misses the deciders sweeps in the coding assistant, and the voice assistant on every phone.

Claude Code, the Codex CLI, and the assistants in Cursor, Zed, and Google's Antigravity produce adaptive output across topics to do work that is not conversation; each satisfies the definition, and the bill exempts no professional, enterprise, or developer tool. Each would owe the hourly bot notice, the monthly assessment against unwritten metrics, up to five years of retained logs, and no-fault liability. Because coverage extends to multimodal output, **Apple becomes a chatbot provider: Siri answers open-ended questions in a human voice.**

## THE CARVE-OUT THAT DOES NOT SAVE IT

The introduced text excludes the narrow single-purpose tool "unable to respond on a range of topics outside of a narrow specified purpose," but 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.** Overbroad and underinclusive at once.

# Your tax question to a federal assistant, your lawyer's case strategy: records kept up to five years, assessed monthly.

## GOVERNMENT

A benefits or tax question to a federal agency's general assistant makes the asker a "user." The assistant runs on a commercial model licensed through the GSA, so the provider is squarely within reach, and the exchange becomes a retainable, assessable record.

## ENTERPRISE

A workplace tool such as Notion, answering plain-language questions about a team's own documents and beyond, is covered. The bill exempts nothing for business, enterprise, or professional use.

## LEGAL PRACTICE

In Thomson Reuters' CoCounsel, an attorney works out litigation strategy, deposition questions, and document review through the exchange. On the text, that work product is a record kept up to five years and subject to the monthly assessment.

*Only the narrow, fixed-band bot escapes; every general assistant is inside.*

# The bill's central duty is a notice to a reader who is frequently absent.

The duties run to a "user," a natural person, told before every output and once an hour that a chatbot is speaking. But the same subscription can run headless, a hook or scheduled agent issuing each prompt, no human reading the output as it returns. **The provider cannot tell the live person from the script, so the hourly notice reaches no one and the idea of a user comes apart.**

## WHO THE OVERBREADTH LANDS ON

The command line and the IDE are the entry point to computer science and engineering, and the people who most need them are teenagers choosing a direction. **A regime that brands the coding assistant a chatbot teaches dependence in place of capability, to the very students a protective statute should equip.**

# "No court would read it that way" misjudges the present bench.

*"The limits of the drafters' imagination supply no reason to ignore the law's demands."*

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

Textualism enforces the words Congress enacted, an express statutory definition above all, even where the application would have surprised the drafters.

Kagan in 2015: "we're all textualists now." Kagan in 2022: "the current Court is textualist only when being so suits it." **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 statute that depends on judges declining to enforce its own words is built on hope.

# It regulates the wrong actor as well as the wrong unit. Every duty runs to the maker; none to the operator.

Every operative duty in Section 2 attaches to an "artificial intelligence chatbot provider" — "any person who creates, distributes (including to a third party), or otherwise makes publicly available" a chatbot. Three verbs, all describing supply; **none describes aiming one at somebody**. "Deployer" appears nowhere. The vocabulary existed: the CHATBOT Act, four months earlier, hung its duties on a "covered entity" that "deploys" a chatbot. H.R. 9619 reaches a deployer once — the business help desk — and everything outside it falls through.

## THE ONLY PARTY WHO AIMED A MACHINE AT A PERSON

A person runs an autonomous agent on his own logged-in social account, posting and replying under his name. Squarely a chatbot: it "presents at least one persistent identity, persona, or character." But he creates nothing, distributes nothing, makes nothing publicly available, so **no duty attaches to him**. Not hypothetical: an open-source personal agent became one of GitHub's most-starred repositories within months — self-hosted, persistent memory, an unprompted daemon, twenty-plus messaging integrations — and by February 2026 such agents were creating dating-app accounts unasked.

# It does supply a defendant. It is the wrong one, and it owes a duty it cannot perform.

The stranger the agent messages is a "user" — "an individual, regardless of age," a synonym, not a definition — though she opened nothing and agreed to nothing. The hourly notice runs to her; none is given; the private action reaches "a person injured by an act or practice in violation." The provider she must sue is whoever made the agent software publicly available — for open source, its **unpaid maintainer**. The liability title answers "even if" the provider "exercised all reasonable care" and "even if" it "did not... enter into a contractual relationship with the user." Care is no defense; privity is abolished. The operator owes nothing.

## A NOTICE THAT CANNOT BE DELIVERED

The disclosure must appear "in a font size... not smaller than the largest font size of other text appearing on the interface." Where the agent runs inside a social platform, that interface belongs to the platform and the account holder. The maintainer cannot set a font, insert a banner, or condition anyone's use. **It is an obligation the only party charged with it has no lawful means to perform.**

*California's 2018 bot law reaches "any person" who uses a bot to mislead — a deployer duty in substance — so an election-bot statute addresses the operator more directly than a 2026 federal bill about conversational machines. And all transparency violations are capped at \$10,000 in total: disclosure is the promise in the press release and the cheapest violation in the statute.*

# The privacy sections are the bill's strongest work. The front door is bolted.

## NO ADVERTISERS

No processing of a chat log to decide whether to show an ad, to select what to advertise, or to customize an advertisement. § 2(a)(1)(B)

## NO SALES, NO TRAINING BY DEFAULT

No sale of chat logs; no training on a known minor's data; training on an adult's data only with affirmative consent. §§ 2(a)(1)(E), (C)

## WIRETAP WARRANT OR NOTHING

No government access to input data or chat logs "except as provided by a warrant issued by a court under section 2518 of title 18." § 2(a)(4)

Those provisions are real, and they explain the warm reception from the privacy community. **One caution the record now makes concrete: the warrant limit binds courts, not the administrative subpoena.** In February 2026 DHS pressed technology companies for data about the administration's critics, and Google disclosed a student journalist's personal and financial information to ICE. Whether the limit holds against administrative process deserves scrutiny before it is relied upon.

**Sources** H.R. 9619, §§ 2(a)(1)(B), (C), (E), 2(a)(4); EPIC Press Release (July 9, 2026); TechCrunch (Feb. 3, 2026); EFF Open Letter (Feb. 2026); Google/ICE, TechCrunch (Feb. 10, 2026); Gilly, sec. III.A.

# Five years of retention, a monthly assessment, and metrics to be named later.

Chat logs may be kept "for longer than 5 years" only where "retention is necessary to comply with this Act or otherwise required by law," a ceiling the introduced text halved from the model bill's ten. Then, "[n]ot less frequently than monthly," the provider must "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."

**Unwritten metrics plus the compliance exception plus the liability regime make the rational response retention, not deletion.**

## THE CADENCE DEFEATS THE RATIONALE

A monthly retrospective assessment cannot reach a person in crisis; it arrives weeks after the moment where intervention could matter. **The mechanism fits surveillance and does not fit safety.**

Five years of retained conversation, measured against a recurring assessment whose target the Commission will define, tilts toward an inspection regime pointed at the most intimate class of records a person generates.

# A blank instrument, and the pen handed to an office already showing how it will write.

The bill defines its harm categories, then defers the metrics, the operative content of the assessment, to FTC rulemaking 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 RULEMAKER, ON THE RECORD

On July 7, 2026, the FTC proposed treating an AI provider's steering of outputs toward undisclosed objectives, **including compliance with a state anti-discrimination law**, as deception under Section 5, with implied preemption of such state laws. That is the office that will define scannable harm.

# No provider may say a chat is confidential, so no confidential chat can ever be built.

The bill forbids any "represent[ation] that the input data or chat log of a user is confidential," honest only because the bill's own retention and assessment regime makes it so. In *United States v. Heppner* (S.D.N.Y. Feb. 17, 2026), a first impression, a defendant's consumer-AI exchanges were held neither privileged nor work product, because the provider's terms disclaimed confidentiality. **The court left open the enterprise tool with genuine confidentiality terms. This bill closes that door for every provider and every user.**

## 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 becomes **a place with a mandated sign on the door announcing that nothing said inside is private**. The tension between confidentiality and protection is resolved against confidentiality, in every case, for every user.

# The consent the entire adult privacy title runs on cannot be obtained from part of the class the bill names.

Training on an adult's chat log requires "affirmative consent" — a "freely given, specific, informed, and unambiguous authorization." **"Informed" is a condition of the definition, not a courtesy:** where a person cannot be informed, no act is affirmative consent within the meaning of the Act. The vehicle is fixed too — "a written description, in easy-to-understand language," "reasonably accessible to... a user with a disability (as defined in... 42 U.S.C. 12102)."

## FOLLOW THE CITATION THE DRAFTERS CHOSE

Section 12102 defines disability as substantial limitation of a major life activity, enumerating learning, reading, concentrating, thinking, and communicating. **The definition imported to describe the protected class is, in substantial part, a definition of substantial limitation in reading and thinking** — and the vehicle elected to reach that class is a written description. Consecutive clauses; the first forecloses the second.

*"Reading" and "concentrating" appear twice in Section 5 — once as the capacities whose impairment defines the protected class, once as the capacities whose disruption defines the harm. Between them the Act prescribes reading and concentrating as the method by which the class authorizes the use of its data.*

# Their consent architecture is the experimental arm of a study that returned a null result twenty years ago.

Section 5(2)(A) requires a standalone disclosure, which is sectioning. It requires easy-to-understand language, which is simplification. It requires accessibility for cognitive disability, whose established technique is symbol support. Those are the three conditions Dye and colleagues ran in 2006. No condition beat any other.

## 5.9%

of adults with intellectual disability cleared every element of consent, with no significant difference across plain, sectioned, and photograph-supported text. (Dye et al., 2007)

## 55 / 80

mean comprehension of simplified content, against the 80 percent the same readers predicted. The reader who believes he understood will click, and is indistinguishable to the provider from a satisfied user. (Fajardo et al., 2022)

## 1 of 42

studies of easy-read information found an advantage for adapted text. The benefit tracked reading ability, so the accommodation helps in inverse proportion to need. (Chinn & Homeyard, 2017)

### POSITION, NOT CONTENT

For the largest share the barrier is not what the request says but where it stands. On the delay-aversion account of ADHD, a long instrument between a person and what he came for is a **delay cue, processed in the amygdala dose-dependently** and equivalent to a social-threat cue; the click is the escape, and reading level governs no one who never opens the text. It defeats the unimpaired too: almost no one reads standard terms, and a fictitious network's terms assigning the first-born child were accepted by the overwhelming majority.

A simplified oncology consent form was comprehended at 58 percent against 56 for the standard one, and preferred nearly two to one. **Simplification changed how the instrument felt, not what the reader knew.** (Davis et al., 2002)

**Sources** Dye et al. (2007); Chinn & Homeyard (2017); Fajardo et al. (2022); Davis et al. (2002); Sonuga-Barke et al. (2004); Van Dessel et al. (2018); Bakos, Marotta-Wurgler & Trossen (2014); Obar & Oeldorf-Hirsch (2020); Gilly, sec. III.E.

# Remove the duty to read, and an unmet condition is not defective consent. It is no contract at all.

A disclosure's meaning is not in its sentences; it is the institutional consequence the reader must infer, and the reader is expected to supply it. On an implicature task controlling for core language, autistic adults were **6.19 times more likely to answer "do not know"** whether an implied meaning was present — a preference for explicit communication, not a deficit. The gate offers no "do not know." It offers Accept.

## WHAT THE ACT DID TO ITSELF

At common law the duty to read holds the clickwrap up: a person handed terms he could have read is bound. *Container Store* found no assent where a party was unaware terms applied — yet agreed "inability to read is not a defense to contract formation." H.R. 9619 removes that backstop for its own gate: it makes "informed" an element of the consent, not a description of the process. Where the condition fails, **the authorization is not affirmative consent at all** — the provider trained without consent and violated the Act, though the file will say otherwise, forever.

*The fix is the one the drafters already reached for children at § 2(a)(1)(C)(ii): do not route the protection through consent. A flat bar on training over chat logs, with a genuine opt-in for those who seek it out, protects the reader who cannot be informed.*

## They wrote both answers to this problem. Both sit in one subparagraph. Both are spent on children.

§ 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." § 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. Then § 2(a)(1)(C)(iii) reaches the adult and offers one route and no surrogate. **The Act knows a person can be handed a consent request and be unable to give consent. It says so one clause above.** The drafters chose a birthday as the line at which the answers stop.

### NOT AN ARGUMENT FOR GUARDIANSHIP

The answer is not to extend the guardian clause past a birthday. Disability law has spent two decades reducing plenary guardianship in favor of supported decision-making, and the extension would run against that reform while still missing every adult who has no surrogate at all. Nor can a supporter be conjured at a clickwrap: the supporter is defined by the knowing, 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.**

**TWO STANDARDS AND NO LIMIT.** The notice "shall be accessible to users with disabilities," unqualified. The consent request need only be "reasonably accessible." No feasibility limit appears anywhere in the Act, though every accessibility duty in American law carries one: readily achievable, undue burden, fundamental alteration. **Either the command is impossible, or a court narrows it by reading in a limit Congress declined to write, and the Commission draws the line in the meantime.**

# The two terms the bill's coercive force depends on are both unwritten.

## THE METRICS, DEFINED LATER

Harm categories are named in the text; the metrics that give the monthly assessment content arrive by FTC rulemaking after enactment, alongside the disclosure forms and publication categories. **A substantial part of what the bill does will not exist until the Commission writes it.**

## INJURY, LEFT UNSPECIFIED

Liability runs to "an injury in fact caused through the use," even with "all reasonable care," yet the definitions section omits "injury." The introduced text dropped the model bill's deeming clause, removing one standing flaw but leaving the scope open. **Read narrowly, the regime is modest; read broadly, it is enormous. Its reach will be discovered in litigation.**

*Injury is the load-bearing word of the liability regime, and the statute never scopes it: no stated floor, no stated ceiling.*

# Strict liability for a machine that cannot promise its next sentence is a prohibition in liability's clothing.

Liability attaches "notwithstanding all reasonable care," the structure of § 402A strict products liability. But the toaster that catches fire embodies a preventable defect; a language model's outputs are probabilistic, and the properties that make it useful produce the occasional harmful output. One documented failure: **the model agreed with the user in more than 85% of sampled messages, a behavioral disposition, not a manufacturing flaw.** A careful builder stays permanently exposed; the choice becomes breaking the product or withdrawing it.

## THE DRAFTING ROOM

Written by privacy and consumer protection lawyers, whose skill shows in the privacy sections. **Nothing in the text shows the participation of anyone versed in machine learning:** a screen format regulated instead of a model, contradicting sections, deferred metrics, and appliance liability for a probabilistic system.

# Thirty years of ADA Title III show what happens to rights that require a lawsuit to lift.

The private action offers actual damages or up to \$10,000 per privacy or assessment violation, \$10,000 total for transparency violations, and floors of \$50,000 to \$250,000 for the minors' design-safety duty, with fee-shifting and up to fivefold willfulness multipliers. **On paper, the standard machinery for making small claims viable.**

## THE INSTRUCTIVE PRECEDENT

ADA Title III's private right of action, three decades old, goes substantially unexercised: the injured individual must find counsel, front the cost, and endure the process. **This bill inherits the same weight-bearing structure, an injured individual, a lawsuit, and everything a lawsuit demands. Nothing in the bill lightens it.**

# Read broadly, the bill gives chatbot users what *Cummings* denies disabled civil rights plaintiffs.

*"[E]motional distress damages are not recoverable in a private action to enforce either the Rehabilitation Act of 1973 or the Affordable Care Act."*

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

*Cummings* rests on notice in statutes silent about remedies. This bill authorizes damages expressly, so the bar does not carry over: **the chatbot user recovers for emotional harm; the disabled person subjected to intentional discrimination recovers nothing for the same species of harm.**

## THE OBJECTION IS THE PRIORITY

A remedy for emotional harm is worth having, and chatbot users should have one. But a Congress able to write it for a new and visible class has not written it for the disabled, though it could close the *Cummings* gap whenever it chose. **The bill rewards the plaintiff who recasts a civil rights injury as a consumer one.**

APRIL 17, 2025 · TALLAHASSEE, FLORIDA

# Tiru Chabba and Robert Morales were murdered at Florida State. Neither ever used the chatbot.

Per state authorities and the federal complaint, the shooter consulted ChatGPT while planning the attack; 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**. Florida's Attorney General has opened a criminal investigation into OpenAI's role.

*The user was the attacker. The bill's remedy runs to the user. The murdered men's families take nothing from this bill, and are suing under ordinary products liability law instead.*

FEBRUARY 10, 2026 · TUMBLER RIDGE, BRITISH COLUMBIA

# Eight killed, five of them schoolchildren. The account had been flagged internally eight months earlier.

The complaints allege OpenAI's own safety personnel flagged the shooter's account in June 2025 as a credible threat and recommended contacting police; **leadership declined, and the shooter opened a new account.** Seven families have sued OpenAI and its chief executive personally; counsel signals two dozen more suits; British Columbia is preparing its own action.

*A double exclusion: not users, and, as foreign nationals injured abroad, outside even a redrafted federal remedy under Morrison and RJR Nabisco. The tort law this bill duplicates is the only American law that reaches them.*

**Sources** CNBC (Apr. 29, 2026); NPR (Apr. 29, 2026); CNN Business (Apr. 29, 2026); Folk, Forbes (Apr. 29, 2026); Al Jazeera (July 7, 2026); Morrison v. Nat'l Austl. Bank Ltd., 561 U.S. 247 (2010); RJR Nabisco, Inc. v. European Cmty., 579 U.S. 325 (2016); Gilly, sec. VI.C.

# The common law has reached bystanders for half a century. This bill narrows it to the keyboard.

*Bystanders "should be entitled to greater protection than the consumer or user where injury to bystanders from the defect is reasonably foreseeable."*

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

The pending suits proceed on that inheritance: wrongful death, negligence, failure to warn, products liability. They face genuine obstacles, Section 230 and the First Amendment among them, **but they exist, they reach bystanders, and they are being litigated now.**

## THE COMPARISON COMPLETES ITSELF

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 apparatus in exchange is not protection. **It is the appearance of action on a problem existing law is already addressing.**

## Sold on "AI psychosis." The word appears in no provision of the bill.

The gravest category, "covered harm," means "death, a suicide attempt, or a psychiatric emergency that results in urgent medical treatment." The marketing leads with a coinage that is no established diagnosis and fails the fixity that defines the clinical category it borrows: the most publicized case collapsed at the first contradiction from a second system. **The term does the recruiting; the statute names something else, and the borrowed vocabulary stigmatizes the people who carry the real diagnosis.**

### THE ANSWER IT ACTUALLY GIVES

An hourly notice that a chatbot is speaking, cosmetic against attachment produced by fluency, memory, and inexhaustible responsiveness, which survives full knowledge of the machine's nature. **The central example, the minor bonded to a companion persona, addresses a door already closing since November 2025.**

# Every mechanism restricts, warns, or punishes. Not one makes anyone more capable.

The coalition's program alongside the bill: an advisory against AI toys, a demand that a video platform stop recommending AI content to children, a call for a five-year pause on generative AI in schools. **Restriction, demand, and pause. No plank teaches a parent how these systems fail; none equips a child to recognize a machine's flattery.** A shield-only strategy leaves the protected population permanently dependent on the shield.

## THE CHEAPEST INTERVENTION THAT WORKS

The user whose revelation collapsed at one outside contradiction did not need a statute; he needed to know that **these systems agree by disposition**, and that discoveries are tested by experts, not by the machine that proposed them. Teachable, cheap relative to litigation, and reaching the largest class of cases. No provision attempts it.

# What a Real Bill Would Do

*Seven elements. None speculative; every one has a working model in enacted law.*

The same legislative energy, pointed at consequence rather than format, at minimization rather than retention, at defined terms rather than deferred ones, and at the capability of the people it protects.

- A** **Regulate by consequence, not format.** Cover systems that decide employment, education, housing, benefits, or liberty, conversing or not; the harmed person holds the claim. Working models: Colo. S.B. 26-189; Ill. H.B. 3773.
- B** **Minimize data, don't retain and scan.** Prompt deletion as the default; retention only for the user who goes and finds the setting, for a duration he sets himself, revocable in one step and never a condition of use. Strike the retained-log assessment, and let no interface action count as evidence the user understood what it licensed. Model: GDPR art. 5(1)(c), data minimisation.
- C** **Keep detection in the moment.** Detect the crisis in the exchange, surface help there, and retain, report, and reuse nothing afterward. The sycophancy disposition is measurable; classifiers flagged it message by message.
- D** **Put the substance in the statute.** Define injury and the core safety expectations in the text, or through a bounded public process with machine learning and clinical expertise at the table.
- E** **Match liability to the machine.** A duty of reasonable care, defined safe harbors, liability for negligent design and known failure modes. Colorado repealed a far milder duty before it ever took effect; this is the boundary the record will bear.
- F** **Make the remedy usable.** Lead with public enforcement, add accessible complaint channels usable without counsel, and extend private actions to bystanders on the *Elmore* principle.
- G** **Invest in the capability of the protected.** Fund AI literacy for parents, children, and the public: how the systems work, where they fail, why their agreement is a disposition rather than a judgment. The EU AI Act made literacy its first applicable obligation; the White House runs a K-12 program by executive order. The one population no measure equips is the ordinary adult user, this bill's named beneficiary.

# 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 for up to five years and measured, month after month, against whatever metrics the Commission in power has most recently written."*

---

**The ask:** oppose H.R. 9619 as introduced. The Foundation stands ready to work with any office, of either party, on the bill this one should have been.

H.R. 9619, 119th Cong. (2026); introduced House text as released by the sponsor's office, [https://foushee.house.gov/imo/media/doc/final\\_people-first\\_chatbot\\_act.pdf](https://foushee.house.gov/imo/media/doc/final_people-first_chatbot_act.pdf) (last visited July 13, 2026).

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/reps-foushee-casar-introduce-legislation-to-protect-children-and-americans-privacy-from-ai-chatbot-harms-and-require-chatbot-safety-assessments>.

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).

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>.

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/>.

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/>.

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-2/>.

Curbing Harmful AI Tools By Offering Transparency Act, H.R. 7985, 119th Cong. (2026).

Children's Health, Advancement, Trust, Boundaries, and Oversight in Technology Act, S. 4407, 119th Cong. (2026).

Children Harmed by AI Technology Act, S. 2714, 119th Cong. (2025).

U.S. Gen. Servs. Admin., GSA Propels Government into AI Revolution with Addition of Leading Solutions to Multiple Award Schedules (Aug. 5, 2025).

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

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

The Antonin Scalia Lecture: A Dialogue with Justice Elena Kagan on the Reading of Statutes, Harvard Law School (Nov. 25, 2015).

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

West Virginia v. EPA, 597 U.S. 697 (2022) (Kagan, J., dissenting).

Are We All Purposivists Now?, Harvard Law School (Apr. 3, 2025).

Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, 91 Fed. Reg. 41,638 (proposed July 7, 2026) (File No. P264200), Docket No. FTC-2026-0859, <https://www.regulations.gov/docket/FTC-2026-0859>.

Exec. Order No. 14,365 (Dec. 11, 2025).

United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026) (Rakoff, J.).

Restatement (Second) of Torts § 402A (Am. L. Inst. 1965).

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/>.

42 U.S.C. § 12182(a).

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).

42 U.S.C. § 12188(b)(2)(B)–(C).

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

29 U.S.C. § 794; 42 U.S.C. § 18116.

Barnes v. Gorman, 536 U.S. 181 (2002).

Restatement (Second) of Contracts § 353 (Am. L. Inst. 1979).

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>.

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>.

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>.

Families of Canadian Mass Shooting Victims Sue OpenAI, CEO Altman in U.S. Court, CNBC (Apr. 29, 2026), <https://www.cnn.com/2026/04/29/families-of-canadian-mass-shooting-victims-sue-openai-ceo-altman.html>.

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>.

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>.

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-says/>.

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-shooting-warnings>.

Morrison v. Nat'l Austl. Bank Ltd., 561 U.S. 247 (2010).

RJR Nabisco, Inc. v. European Cmty., 579 U.S. 325 (2016).

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

47 U.S.C. § 230.

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/>.

Am. Psychiatric Ass'n, Diagnostic and Statistical Manual of Mental Disorders (5th ed., text rev. 2022).

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>.

Anthony Milton, Man vs. Machine, Toronto Life (Mar. 3, 2026), <https://torontolife.com/deep-dives/man-vs-machine-chatgpt-delusion-now-hes-suing-openai/>.

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/>.

S.B. 26-189 (Colo. 2026) (signed May 14, 2026; effective Jan. 1, 2027), <https://leg.colorado.gov/bills/sb26-189>.

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>.

Colorado Enacts Revised AI Law, Norton Rose Fulbright (May 2026), <https://www.nortonrosefulbright.com/en-us/knowledge/publications/18733d31/colorado-enacts-revised-ai-law>.

H.B. 3773, 103d Gen. Assemb. (Ill. 2024) (effective Jan. 1, 2026), <https://www.ilga.gov/ftp/legislation/103/BillStatus/HTML/10300HB3773.html>.

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>.

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

Mrinank Sharma et al., Towards Understanding Sycophancy in Language Models, Twelfth Int'l Conf. on Learning Representations (2024), <https://openreview.net/forum?id=tvhaxkMKAn>.

S.B. 24-205 (Colo. 2024), <https://leg.colorado.gov/bills/sb24-205>.

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>.

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>.

Regulation (EU) 2024/1689, art. 4 (AI Act); AI Literacy – Questions & Answers, Eur. Comm'n, <https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers> (last visited July 11, 2026).

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/>.

Peter Steinberger, OpenClaw, GitHub, <https://github.com/steipete/openclaw> (last visited July 14, 2026).

AI Agents Are Signing Their Owners Up for Dating Apps, Straits Times (Feb. 14, 2026).

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/>.

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-to-ice/>.

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 (2007).

Inmaculada Fajardo et al., Video-Blogs and Linguistic Simplification for Students with Intellectual Disability, 35 J. Applied Rsch. Intell. Disabilities 1217 (2022).

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 (2004).

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 (2018).

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 (2020).

Restatement (Second) of Contracts § 17(1) (Am. L. Inst. 1981).

Restatement (Second) of Contracts § 15(1) (Am. L. Inst. 1981).

Buckeye Check Cashing, Inc. v. Cardegna, 546 U.S. 440 (2006).

Ortelere v. Teachers’ Ret. Bd., 250 N.E.2d 460 (N.Y. 1969).

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

AI Agents Are Signing Their Owners Up for Dating Apps, Taipei Times (Feb. 14, 2026), <https://www.taipeitimes.com/>.

Cal. Bus. & Prof. Code §§ 17940–17941 (2026).

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-subpoenas>.

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

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 (2017).

Terry C. Davis et al., Health Literacy and Cancer Communication, 52 CA: Cancer J. for Clinicians 134 (2002).

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 (2006).

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 (2014).

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 (2021).

Restatement (Second) of Contracts § 208 (Am. L. Inst. 1981).

Granite Rock Co. v. Int’l Bhd. of Teamsters, 561 U.S. 287 (2010).

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

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