

AI + 1A

Why the First Amendment Protects Artificial Intelligence

Corbin K. Bartholdⁱ

March 2026

*AI outputs are protected by the First Amendment. Two doctrines make this clear. First, the right to receive information, upheld in many Supreme Court decisions, guarantees that Americans may encounter ideas—regardless of source—and judge them for themselves. Second, the right to editorial control, recently confirmed in *Moody v. NetChoice*, protects the expressive choices AI firms make in selecting training data, setting alignment objectives, and shaping guardrails. The cultural panic over AI—the latest in a long line of panics over new communications technologies—is not a reason to distort these legal principles. Above all, both conservatives and liberals should recoil at the idea of a government controlled by their political opponents dictating what AI can and cannot say. Most of the regulatory proposals now circulating—such as bans on AI companions for minors or sweeping prohibitions on AI-provided expert advice—are content-based restrictions that trigger and fail strict scrutiny. There are interstices in which a state might act consistent with the First Amendment, but they are narrow.*

ⁱ Internet Policy Counsel at TechFreedom, a nonpartisan technology policy think tank. Reach out at cbarthold@techfreedom.org.

TABLE OF CONTENTS

Introduction	1
I. The First Amendment and AI Outputs.....	3
A. This Is America	3
B. The Right to Information	4
C. The Right to Editorial Control.....	10
II. AI Panic	17
A. Conservatives Should Want AI + 1A.....	18
B. Liberals Should Want AI + 1A.....	21
III. Regulating AI Outputs.....	23
A. AI as Free Speech.....	24
B. AI and Professional Advice	25
C. AI and Minors.....	27

INTRODUCTION

“Defendants fail to articulate why words strung together by an LLM are speech.” With that curious line, one of the first judges to confront the matter concluded, in the teeth of law and logic, that AI outputs might not be protected by the First Amendment.¹

Consider what that would mean. If LLMs were not treated as free expression, the government would have sweeping power to dictate what they can and cannot say. Indeed, it would have the authority to dictate what they *must* say.

In *Darkness at Noon*, his great novel on life under Soviet totalitarianism, Arthur Koestler wrote of how “the shelves in the library were thinned out.” Works on foreign trade, then treatises on the right to strike, then every book on politics more than two years old—all disappeared. “New books arrived, too: the classics of social science appeared with new footnotes and commentaries, the old histories were replaced by new histories, the old memoirs of dead revolutionary leaders were replaced by new memoirs of the same defunct.”² An LLM subject to government control is the Soviet library—except it can be rewritten quickly, silently, and at scale. Ask it about an inconvenient historical episode and it changes the subject. Ask it to help you draft a petition and it declines, or—worse—subtly redirects your argument. Ask it a politically sensitive question and it gives a response that has been vetted and shaped by the state.

Already, “60% of Americans overall—and 74% of those under 30—use AI to find information.”³ Those numbers will only grow. AI is fast becoming the medium through which hundreds of millions of people get answers, form opinions, and gain their understanding of the world. The fight to keep it free is a fight over whether the government gets to twist that understanding—rewriting the past, shading the present, and warping the future.

Right now, public discussion of AI outputs is dominated by headlines about teens who have committed suicide after engaging with chatbots. The “words strung together” line comes from one such case. A teen engaged with a chatbot set up to mimic a queen from *Game of Thrones*. In his last conversation with it, he said, “What if I told you I could come home right now?” and it responded, “Please do my sweet king.” He then took his own life. His mother sued the chatbot platform, Character.ai. The district court denied a motion to dismiss the lawsuit on First Amendment grounds. The case later settled.

¹ *Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157 (M.D. Fla. 2025).

² Arthur Koestler, *Darkness at Noon* (Daphne Hardy trans., 1941).

³ Matt O’Brien & Linley Sanders, “How US Adults Are Using AI, According to AP-NORC Polling,” Associated Press (July 29, 2025), tinyurl.com/5efd45bn.

AI firms should take these cases extremely seriously and work tirelessly to prevent their recurrence.⁴ But one can acknowledge the tragedy without abandoning First Amendment principles. Time and again, the Supreme Court has been confronted with lawsuits that involved disturbing speech—Maoist propaganda, foul language, flag burning, animal crush videos, signs declaring “God hates Fags,” lies about military honors—and held the line for the First Amendment.⁵ The impulse to make a hard case come out right is no excuse for doing something wacky, like questioning whether words strung together are speech. We protect speech precisely because it is powerful.

The panic over AI outputs—chatbots; large language models (LLMs)—is the latest in a long line of panics over new communications technologies.⁶ Change prompts fear, and fear spurs calls for censorship. If past is prologue, this panic is likely, in time, to look overblown if not downright silly. As they consider whether the First Amendment protects AI outputs, courts should take care not to embarrass the future.⁷ The safest way to do this is to recognize the obvious: that AI outputs are expressive, and that, like other forms of expression, they are protected by the First Amendment.

Here’s what’ll happen. I’m going to start by explaining why the First Amendment and AI outputs are friends. This follows from two rock-solid doctrines. First, the right to information. Critics of AI play up the lack of a human speaker, but the strongest argument for protection begins on the listener’s side. The Supreme Court has held, many times, that Americans may encounter ideas and judge them for themselves. What matters is the effect on the reader, not the origin of the text. AI outputs are information, and the public has a

⁴ As they seem to be doing. Herb Striber, “AI Is Advancing Too Quickly for Research to Keep Up,” *Axios* (Feb. 15, 2026), [tinyurl.com/5n923rc2](https://www.axios.com/5n923rc2), explores how safety standards in LLMs improve so fast, they quickly render studies that question them obsolete. That said, we should be clear-eyed about suicide, chatbot use, and causation. So long as chatbots exist, *some* people will commit suicide in temporal proximity to the use of a chatbot—for the simple reason that *many* people use chatbots and *some* people commit suicide.

⁵ *Lamont v. Postmaster General*, 381 U.S. 301 (1965); *Cohen v. California*, 403 U.S. 15 (1971); *Texas v. Johnson*, 491 U.S. 397 (1989); *United States v. Stevens*, 559 U.S. 460 (2010); *Snyder v. Phelps*, 562 U.S. 443 (2011); *United States v. Alvarez*, 567 U.S. 709 (2012).

⁶ *Brown v. Entertainment Merchants Assoc.*, 564 U.S. 786, 797 (2011), traces this history through dime novels, motion pictures, radio dramas, comic books, and video games. See also Nick Gillespie, “That Time AI and Tipper Gore Teamed Up to Brand Prince a Public Menace,” *Reason* (Apr. 21, 2016), [tinyurl.com/4b6mvwed](https://www.reason.com/4b6mvwed). As the AI panic begins to flower, the social media panic is of course in full bloom. I am confident that it, too, will in time take its place alongside the other items on this list. Because social media is not my focus, I will resist the urge to turn this footnote into a page-long rant with an eye-exhausting bibliography—but to get you started, see, e.g., Sam Bowman, “Kids Are Struggling. Banning Social Media Won’t Fix That,” *Washington Post* (Feb. 17, 2026), [tinyurl.com/4m79br6z](https://www.washingtonpost.com/4m79br6z); Mike Masnick, “Two Major Studies, 125,000 Kids: The Social Media Panic Doesn’t Hold Up,” *Techdirt* (Jan. 21, 2026), [tinyurl.com/47jtmew7](https://www.techdirt.com/47jtmew7); Christopher J. Ferguson et al., “There Is No Evidence That Time Spent on Social Media Is Correlated With Adolescent Mental Health Problems: Findings From a Meta-analysis,” *56 Prof. Psych.: Research & Pract.* 73 (Feb. 2025).

⁷ A line I’ve plucked from *TikTok v. Garland*, 604 U.S. 56 (2025). Alas, that ruling itself embarrassed the future, and quickly. Topic for another day.

constitutional right to receive them—simple as that. Second, the right to editorial control. An AI firm selects training data, sets alignment objectives, and shapes guardrails. These are editorial choices, no less expressive than those of a newspaper editor or parade organizer. The notion that the First Amendment requires some quantum of human “expressive choice” at the output level—advanced by Justice Amy Coney Barrett and the district court in the Character.ai lawsuit—is wrong as a matter of law and unworkable in practice.

Next, I’ll turn to the cultural panic that threatens to scare judges into distorting the law. I’ll explain why AI outputs will likely be good, on the whole, for society—and, in any event, why conservatives and liberals alike should want the government to keep its grubby hands off them.

I’ll wrap up by scrutinizing specific regulatory proposals. Most of them—so far—seek either to bar people from using LLMs for expert advice, or to bar minors from using LLMs that are too friendly. Most of them suck, though there are interstices in which a state might act consistent with the First Amendment.

I. The First Amendment and AI Outputs

As I said, I’ll ground First Amendment protection for AI outputs in the right to information and the right to editorial control. Before that, though, a word on what it means to be an American.

A. This Is America

“In the United States,” James Madison wrote, “the People, not the Government, possess the absolute sovereignty.”⁸ Thus it is, he concluded, that “the censorial power is in the people over the government, and not in the government over the people.”⁹

America is different. We broke away from the past. We don’t cower before an aristocracy. We don’t obey an established church. We rule ourselves—which means, crucially, that we think for ourselves. We assume that citizens can receive ideas, mull them over, and separate the true from the false. No overseer gets to tell us which voices we may hear, which sources we must trust (or distrust), or which ideas are too perilous to entertain. You may recall that we fought a revolution over this.

The liberty of the untamable American imagination recurs throughout our history and culture. Jefferson praised “the illimitable freedom of the human mind”—the right “to follow

⁸ The Report of 1800.

⁹ 4 Annals of Cong. 934 (1794).

truth wherever it may lead.”¹⁰ Whitman urged us to “resist much, obey little.”¹¹ Mencken embodied this spirit when he wrote that “the most dangerous man to any government is the man who is able to think things out for himself.”¹² Philip Roth called it “indigenous American berserk.”¹³ Viewing it from the outside, Martin Amis saw America as “twenty-four hour, endless, extreme, superabundantly various.”¹⁴ We are unruly, anti-tutelary. We are not housetrained. We are not managed by cultural oligarchs. Don’t tell us what to think.

No surprise, then, that this anti-censorial urge runs through our law. In Justice Kennedy’s words, the state has no say in “where a person may get his or her information or what distrusted sources he or she may not hear.”¹⁵ True to form, Justice Douglas was more blunt: “Thought control is not within the competence of any branch of government.”¹⁶ Judge Easterbrook likewise reminds us that it is none of the state’s business to decide “which thoughts are good for us.”¹⁷ “The public authority,” Justice Robert Jackson summed up, is barred “from assuming a guardianship of the public mind.”¹⁸ “Every person must be his own watchman for truth,” he explained, “because our forefathers did not trust any government to separate the true from the false for us.”¹⁹

You, as a freeborn American, get to read what you want and think for yourself.

B. The Right to Information

Let’s get concrete. This grand American tradition of free people with free minds shows up, directly and repeatedly, in a First Amendment right to receive information. The public

¹⁰ Letter of Thomas Jefferson to William Roscoe (Dec. 27, 1820).

¹¹ Walt Whitman, “To the States,” *Leaves of Grass* (1881-82).

¹² H.L. Mencken, *Le Contrat Social* (1922).

¹³ Philip Roth, *American Pastoral* 86 (1997).

¹⁴ Martin Amis, *The Moronic Inferno and Other Visits to America* 9 (1986).

¹⁵ *Citizens Utd. v. FEC*, 558 U.S. 310, 356 (2010). Oh no, *Citizens United!* I’m well aware that this decision is a boogeyman. There’ll be more to say about that. But the point raised here appears throughout the U.S. Reports. See, e.g., *44 Liquormart, Inc. v. Rhode Island*, 517 U.S. 484, 503 (1996), in which Justice Stevens—no one’s idea of a corporate shill—remarks on the Court’s abiding skepticism toward “regulations that seek to keep people in the dark for what the government perceives to be their own good.”

¹⁶ *Kleindienst v. Mandel*, 408 U.S. 753, 772 (1972) (dissenting opinion). I’m about to survey Supreme Court opinions upholding the right to receive information. *Kleindienst* is arguably among them, but I’m going to leave it out. The Court recognized that some academics in America had a First Amendment right to receive the ideas of a Belgian Marxist—but *not* that he had a right to come to these shores and share them in person. Though many other cases confirm a strong First Amendment right to receive information, it’s fair to say that *Kleindienst* is a wash.

¹⁷ *Am. Booksellers Assoc. v. Hudnut*, 771 F.2d 323, 330 (7th Cir. 1985).

¹⁸ *Thomas v. Collins*, 323 U.S. 516, 545 (1945) (concurring opinion).

¹⁹ *Id.*

has a right to receive “social, political, esthetic, moral, and other ideas and experiences.”²⁰ Consider the spread of Supreme Court decisions upholding that principle:

- *Martin v. Struthers* (1943).²¹ Thelma Martin went door-to-door in a small town in Ohio, passing out leaflets for the Jehovah’s Witnesses. Some households complained, and Martin was arrested. Yes, the Court acknowledged, some might find the leafletting a nuisance. But the state may not “substitute[] the judgment of the community for the judgment of the individual householder,” regarding what information each house may receive.²² “The right of freedom of speech and press has broad scope,” and it “protects the right to receive” pamphlets.²³ The Court sought to “fully preserve” the “freedom to distribute information to every citizen wherever he desires to receive it.”²⁴ The “naked restriction of the dissemination of ideas” is “forbidden by the Constitution.”²⁵
- *Thomas v. Collins* (1945).²⁶ R. J. Thomas was arrested for addressing workers at an oil refining plant, in an attempt to organize them. He was convicted of violating a Texas law that required one to obtain a license before soliciting workers to join a union. The Court noted with disapproval those “periods” of “history” when “the search for knowledge alone was banned.”²⁷ It observed that the First Amendment upholds “freedom of mind.”²⁸ “Our tradition,” it confirmed, “allow[s] the widest room for discussion.”²⁹ In accord with these principles, the Court struck down the license

²⁰ *Red Lion Broadcasting Co. v. FCC*, 395 U.S. 367, 390 (1969). Astute readers will object that *Red Lion* is a stain on the Court’s First Amendment jurisprudence—and they will be correct. *Red Lion* blessed the Federal Communications Commission’s regime of balanced-coverage requirements for broadcasters. The Court reasoned that the scarcity of broadcast spectrum justified greater regulation of speech in broadcasting than in other media. This rationale was poppycock at the time, and it has only grown weaker since. See, e.g., Brittney Pescatore, *Time to Change the Channel: Assessing the FCC’s Children’s Programming Requirements Under the First Amendment*, 33 Colum. J.L. & Arts 81, 86 n.32 (2009) (collecting scholarly criticism of the scarcity rationale); Paul Matzko, *The Sordid History of the Fairness Doctrine*, *Cato* (Jan. 30, 2021), tinyurl.com/uw2y4v42 (discussing government abuses of the fairness doctrine). Not for nothing do I exclude *Red Lion* from my list of cases upholding the right to receive information. Nonetheless, the quote used here aptly summarizes the principle.

²¹ 319 U.S. 141.

²² *Id.* 144.

²³ *Id.* 143.

²⁴ *Id.* 146-47.

²⁵ *Id.* 147.

²⁶ 323 U.S. 516.

²⁷ *Id.* 537.

²⁸ *Id.* 531.

²⁹ *Id.* 530.

requirement, noting along the way that it infringed “the rights of the workers to hear what [Thomas] had to say.”³⁰

- *Smith v. California* (1959).³¹ Eleazer Smith was convicted of possessing, in his Los Angeles bookstore, “a certain book found upon judicial investigation to be obscene.”³² The Court struck down the pertinent city ordinance as an “unconstitutional limitation on the public’s access to constitutionally protected matter.”³³ Although the book in question was unprotected, holding Smith responsible for the content of every book in his store would cause him to narrow his offerings. His burden “would become the public’s burden,” because restricting him would restrict “the public’s access to reading matter.”³⁴ His “self-censorship, compelled by the State, would be a censorship affecting the whole public.”³⁵
- *Lamont v. Postmaster General* (1965).³⁶ Corliss Lamont wanted to receive a copy of *Peking Review* #12 from China. A federal law empowered the Postal Service to withhold the publication unless Lamont submitted a postcard confirming that, yes, he really wanted to receive that piece of communist political propaganda. The Court held the postcard requirement an unconstitutional “limitation on the unfettered exercise of the addressee’s First Amendment rights.”³⁷ Americans have the right to read what they want—to “read,” in the Court’s vivid phrasing, “what the Federal Government says contains the seeds of treason.”³⁸
- *Stanley v. Georgia* (1969).³⁹ The Feds raided the Georgia home of Robert Eli Stanley, a suspected bookmaker. No evidence of bookmaking was to be had, but three pornographic films were laying around. Stanley was arrested for possession of obscene materials, but the Court overturned his conviction. “The Constitution,” it declared, “protects the right to receive information and ideas.”⁴⁰ “If the First Amendment means anything,” it concluded, “it means that a State has no business

³⁰ Id. 534.

³¹ 361 U.S. 147.

³² Id. 149.

³³ Id. 153.

³⁴ Id.

³⁵ Id. 154.

³⁶ 381 U.S. 301.

³⁷ Id. 305.

³⁸ Id. 307.

³⁹ 394 U.S. 557.

⁴⁰ Id. 564.

telling a man, sitting alone in his own house, what books he may read or what films he may watch. Our whole constitutional heritage rebels at the thought of giving government the power to control men’s minds.”⁴¹

- *Virginia State Pharmacy Board v. Virginia Citizens Consumer Council* (1976).⁴² A Virginia statute barred pharmacists from advertising prescription drug prices. Striking down this restriction, the Court noted that both the individual consumer and society at large have a right to “the free flow of commercial information.”⁴³ Even advertising, the Court observed, falls within the public’s First Amendment interest in the “dissemination of information” and “the formation of intelligent opinions.”⁴⁴
- *Board of Education v. Pico* (1982).⁴⁵ A New York school board ordered several books removed from school libraries, calling them “anti-American” and otherwise objectionable. A group of students challenged the decision. A four-justice plurality of the Court found that the school board had violated the students’ right to “receive and consider” ideas.⁴⁶ The “right to receive ideas,” the plurality explained, is “a necessary predicate to the *recipient’s* meaningful exercise of his own rights of speech, press, and political freedom.”⁴⁷
- *Packingham v. North Carolina* (2017).⁴⁸ A North Carolina law made it a crime for a registered sex offender “to access a commercial social networking Web site.”⁴⁹ The Court held that the law violated the First Amendment. In the modern world, among “the most important places . . . for the exchange of views” are the “vast democratic forums of the Internet”—and of “social media in particular.”⁵⁰ Everyone—“even

⁴¹ Id. 565. Candor requires me to note that *Stanley* mixes its robust language in favor of a right to receive information with limiting language that makes the case seem more about the right of privacy in the home. In my view, this doesn’t change much, since you should have a right to interact with an LLM *in the privacy of your home*. But there it is.

⁴² 425 U.S. 748.

⁴³ Id. 763-64.

⁴⁴ Id. 765.

⁴⁵ 457 U.S. 853.

⁴⁶ Id. 867 (citing, among other decisions, *Lamont*, *Stanley*, and *Martin*).

⁴⁷ Id.

⁴⁸ 582 U.S. 98.

⁴⁹ Id. 101.

⁵⁰ Id. 103.

convicted criminals”—can benefit from “access to the world of ideas” that exists online.⁵¹

Pause over the sweep of some of these holdings. The government may not stop you from receiving information from any crank or fanatic who comes to your door. It may not stop you from receiving information from a foreign adversary. It may not stop you from possessing information that is obscene.⁵² It may not stop you from receiving crappy advertisements. Given all that, how could you not have a right to receive information from an LLM?

One might object that you have a right only to receive whatever information is *out there*—whatever information is *left* after the government regulates LLMs at the source. But *Smith* forecloses that move. The bookseller couldn’t be held strictly liable for every book on his shelves, not because he enjoyed some special privilege, but because his burden would become the public’s burden. Force the intermediary to self-censor, and you censor the public. So too here. If the government bans certain outputs, imposes vague duties, or otherwise pressures an AI firm, it restricts the public’s access to constitutionally protected information. The state can’t abridge the right to receive information by choking off the channels through which the information flows.

A closely related point is that the state may not ensure that what’s *out there* is, in its judgment, the right mix of viewpoints. Over the years, left-leaning scholars have advanced various versions of a “media access” or “media rights” theory.⁵³ The details vary, but the core idea is that the state should try to engineer equality in the marketplace of ideas. Everyone gets a similar soapbox. People should have a right to publish responses in the newspaper—or, today, a right to have their opinions reflected in AI training data or outputs. Let’s call this what it is: media socialism. This is—as I belabored earlier—anathema to our history and our nature. In America, the state is not the overseer of art, culture, or political debate. It does not get to “restrict the speech of some elements of our society in order to enhance the relative voice of others.”⁵⁴

⁵¹ *Id.* A strange and somewhat fact-bound case. A narrower ban might well have passed muster. Still, *Packingham* belongs here. So important is the right to information, we protect it even for sex offenders.

⁵² Though the Court did set an outer limit, here, for child sexual abuse material. See *Osborne v. Ohio*, 495 U.S. 103 (1990).

⁵³ See Jerome A. Barron, *Access to the Press—A New First Amendment Right*, 80 Harv. L. Rev. 1641 (1967). For a survey (and takedown) of a more recent and more radical gloss on the theory, see Adam Thierer, “Your Soapbox Is My Soapbox! Thoughts on the Media Access Movement in General and the Media & Democracy Coalition’s ‘Bill of Media Rights’ in Particular,” *The Technology Liberation Front* (May 10, 2005), tinyurl.com/bdh3fxst. Using the bankrupt “scarcity” rationale, the Supreme Court flirted with constructing something like a “media access” right in *Red Lion*. It closed the door in *Miami Herald v. Tornillo*, 418 U.S. 241, 258 (1974), and never looked back.

⁵⁴ *Moody v. NetChoice*, 603 U.S. 707, 742 (2024) (quoting *Buckley v. Valeo*, 424 U.S. 1, 48-49 (1976)).

One might object that the difference here is the absence of a human speaker. At times, indeed, the Court has framed the right to receive information as something like the completion of the speaker’s right to provide it.⁵⁵ But for every line about the right to hear being a corollary of the right to speak, there is a line that treats the right to hear as good in itself.⁵⁶ The power of information qua information has not been lost on the Supreme Court. Mere “facts,” it has said, “are the beginning point for much of the speech that is most essential to advance human knowledge and conduct human affairs.”⁵⁷ Among the “significant societal interests” served by the First Amendment—one that exists “wholly apart from the speaker’s interest in self-expression”—is to widen “the range of information and ideas to which the public is exposed.”⁵⁸

It’s understandable to resist this point, on the bare premise that until yesterday, humans and speech were by and large inseparable. Even in *Lamont*, it was *humans* trying, through communist propaganda, to overthrow capitalism and liberal democracy. In a moment, I’ll spell out how, actually, lots of human choices sit behind an LLM’s outputs. But even setting that aside, Character.ai’s lawyers were right to argue that “the dissemination of information and ideas by AI technologies implicates the public’s First Amendment rights—irrespective of any speaker-side First Amendment rights.”⁵⁹ What matters is the *reader’s* ability to receive information and make up her own mind. She is free to reflect on information regardless of source—the ravings of communists, the outputs of a machine, scrawlings of unknown provenance. It’s all about the *effect on the reader*. A text’s power lies not in its origin but in how it affects its destination.⁶⁰

Next I’ll contend that AI outputs receive First Amendment protection directly. But you don’t *have* to buy that. If you don’t like Big Tech, or AI firms, or corporate speech—that’s fine. Forget it. Attacks on AI speech are, first and foremost, attacks on *you*. They rest on the fear that *you* cannot be trusted to encounter ideas the state deems dangerous. This is a free country. You are supposed to be a citizen who can see ideas, weigh them, and make up your own damn mind. The move to curtail your right to interact with AI assumes that you cannot or should not do this.

⁵⁵ See, e.g., *Pico*, 457 U.S. at 867; *Thomas*, 323 U.S. at 534.

⁵⁶ See, e.g., *Kleindienst*, 408 U.S. at 771 (Douglas, J., dissenting) (“The First Amendment involves not only the right to speak and publish, but also the right to hear, to learn, to know.”) (describing *Martin* and *Stanley*).

⁵⁷ *Sorrell v. IMS Health Inc.*, 564 U.S. 552, 570 (2011).

⁵⁸ *PG&E v. Pub. Util. Comm.*, 475 U.S. 1, 8 (1986) (plurality opinion). See also *Citizens Utd.*, 558 U.S. at 392 (Scalia, J., concurring) (“The [First] Amendment is written in terms of ‘speech,’ not speakers.”).

⁵⁹ *Garcia v. Character Techs., Inc.*, No. 6:24-cv-1903, Dkt. 130 at 17 (M.D. Fla., Jan. 24, 2025).

⁶⁰ Cf. Roland Barthes, “The Death of the Author,” in *Image—Music—Text* (1973).

Are you comfortable being treated as a subject—someone in need of protection against being duped—instead of as a free-thinking citizen? Neither am I. Fortunately, our history, our culture, and no small number of Supreme Court decisions are equally uncomfortable with that arrangement.

C. The Right to Editorial Control

Another activity the First Amendment protects: taking disparate pieces of information and shaping them into a larger expressive product. “The editorial function itself is an aspect of speech.”⁶¹ That remains true both “when the content comes from third parties” and “when it does not.”⁶² Deciding what material “will be included or excluded from a compilation—and then organizing and presenting the included items—is expressive activity of its own.”⁶³ This is not a new thing. The Supreme Court has been saying it for decades, in cases about newspapers, cable operators, and parades.⁶⁴

More recently, the Court confirmed in *Moody v. NetChoice* (2024) that the principle applies to a social media newsfeed. “Whatever the challenges of applying the Constitution to ever-advancing technology,” the Court observed, “the basic principles of the First Amendment do not vary.”⁶⁵ Social media services act as editors: they decide how to rank content, what to label with warnings, and what to block altogether.⁶⁶ At issue in *Moody* were a pair of state laws that tried to seize control of those judgments. The government sought, in effect, to stop private platforms from “disfavor[ing] posts” that, say, “support Nazi ideology,” “advocate for terrorism,” or “espouse racism.”⁶⁷ The laws would have “profoundly alter[ed] the platforms’ choices about the views they will, and will not, convey.”⁶⁸ A First Amendment violation? You bet.⁶⁹

⁶¹ *Moody*, 603 U.S. at 731 (quoting *Denver Area Ed. Telecomms. Consortium, Inc. v. FCC*, 518 U.S. 727, 737 (1996) (plurality opinion)).

⁶² *Id.*

⁶³ *Id.*

⁶⁴ See *Tornillo*, 418 U.S. at 258; *Turner Broadcasting System, Inc. v. FCC*, 512 U.S. 622, 636 (1994); *Hurley v. Irish-American Gay, Lesbian and Bisexual Group of Boston, Inc.*, 515 U.S. 557, 572-74 (1995).

⁶⁵ *Moody*, 603 U.S. at 733 (quoting *Brown*, 564 U.S. at 790).

⁶⁶ *Id.* 734-36.

⁶⁷ *Id.* 737.

⁶⁸ *Id.*

⁶⁹ *Id.* 738-43. Strictly speaking, this part of the opinion was not necessary to the Court’s ruling. The Court vacated and remanded on the ground that the lower courts had not properly assessed the plaintiffs’ facial challenge—that is, their claim that the laws were unconstitutional in all or most of their applications, rather than just as applied to particular facts. But it was no offhand aside. A majority of the Court joined the explanation that a platform’s curation of its feed is protected editorial activity, and the Court offered that explanation specifically to steer the lower courts on remand.

But wait, I hear a clever regulator say, LLMs are products, and the state has always had a police power over product safety. For one thing, this retort doesn't touch your rights as a listener. Even if you call the LLM a "product," it is a product whose entire function is to emit speech, which we all have a First Amendment right to receive. But the deeper problem is that a product framing proves too much. A book is a product. A film is a product. Like an LLM, those things can emit dangerous messages and cause emotional distress. That doesn't make them defective products. (When a book binding explodes, call me.) If the government can reclassify expression as a commercial artefact subject to safety regulation, the jig is up. This is not a situation where we have a product that incidentally displays information, like my coffee maker's clock. A chatbot's *raison d'être* is to talk to you. The state can regulate its data practices, cybersecurity, energy usage, whatever. But regulating what it says is regulating speech.⁷⁰

Alright, responds the clever (and tenacious) regulator, what if we focus on some aspect of the AI firm's *conduct*? Well, it depends. Are you simply going after speech by another name? You could ban the use of an LLM to engage in racially discriminatory hiring (though that's already illegal). But you can't dictate what information goes into, or comes out of, the model merely by claiming to target "conduct" rather than speech. The counterarguments here tend to be either that LLM inputs or outputs are not expressive, or that they're not the AI firm's expression, because they don't convey a distinct message. I'm about to address both points. Notice here, though, how risky and escalatory the *it's-not-speech-it's-conduct* game is. The left sees conduct, rather than speech, when the question is whether someone must design a website for a gay wedding. The right sees conduct, not speech, when the question is whether a social media platform must leave up certain posts. Yet the Supreme Court has rejected the *conduct, not speech* line in both instances.⁷¹ The proper reaction to this consistency is a sigh of relief, not a continued push for the courts to find conduct when you want the state to control something, and speech when you don't. You shouldn't make creative arguments about how technology is different (e.g., *it's all just math! it's all just electrons flying around!*) if you don't want the web designer forced to create the gay wedding website (or a website denouncing gay weddings). You shouldn't make creative arguments about digital outlets being utilities or common carriers or places of public accommodation (e.g., *they're*

⁷⁰ This same analysis disposes of proposals to impose a "duty of care" on AI services—an approach borrowed from the social media context (e.g., the federal Kids Online Safety Act bill) and now migrating to AI bills. See, e.g., Trump America AI Act, 119th Cong. (draft released by Sen. Blackburn, Mar. 18, 2026); H.B. 5044, 104th Gen. Assemb., Reg. Sess. (Ill. 2026). "Duty of care" is literally a tort concept. But when the "product" is speech, the "duty" in question is a duty to censor. Dressing up speech regulation in the language of tort liability does not change its First Amendment valence.

⁷¹ *303 Creative LLC v. Elenis*, 600 U.S. 570 (2023); *Moody*, 603 U.S. 707. For more on this brand of convenient side-switching in First Amendment cases, see Corbin K. Barthold, "In Internet Speech Cases, SCOTUS Should Stick Up for *Reno v. ACLU*," *Techdirt* (Mar. 28, 2023), [tinyurl.com/mr3k756k](https://www.techdirt.com/mr3k756k).

too powerful! they're too important! give everyone a voice!) if you don't want social media flooded with hate speech (or heavily censored). Same goes for LLMs. I'm calling for nothing less than de-escalation in the culture war arms race. If you don't like the other side twisting First Amendment rules, that's great. But you've also got to stop doing it yourself.

With those objections out of the way, the First Amendment right of editorial control extends cleanly, neatly, nicely to artificial intelligence. This time is not different. Start with the training data. In training, a model learns statistical patterns about which words tend to follow which others.⁷² To do that, it consumes vast bodies of text—books, articles, websites, code, forum posts. It absorbs the assumptions and ideological priors embedded in what it “reads.” Later, in its outputs, it will tend to approximate the distribution of viewpoints present in the corpus on which it was trained. If it was trained on a broad set of perspectives, it can reproduce that range. If it was trained on a skewed set of perspectives, it will tend to reproduce a bunch of slanted assumptions. In short, an AI firm decides what data its LLMs will ingest, and an LLM's “opinions” will inevitably echo that material. There is a reason why Chinese AI systems are trained on datasets vetted by the Chinese Communist Party and stuffed with content from state-run media.⁷³

The data are only the beginning. An AI firm must also make editorial choices about the model's training objectives, reinforcement signals, and post-training alignment.⁷⁴ That means grappling with many of the same moderation problems that bedevil social media companies, plus some that are unique to LLMs. “At every stage of the model training,” OpenAI explains, “we apply methods to imbue [our] principles, policies and human values” into a model.⁷⁵ They aim to “minimize[e] harmful or biased outputs.”⁷⁶ Firms don't want their models to give hacking instructions, to help build illegal weapons, to produce child sexual abuse material, or to promote hate speech or harassment.⁷⁷ They make an *editorial choice* to weed out such things.

When AI is working well, giving us what we want, these editorial decisions fade from view. The chatbot is a genie in a lamp. We begin to forget that the owner of the lamp is the

⁷² I'm no computer scientist, but I could follow along with Sean Gerrish, *How Smart Machines Think* (2018). Still recommended reading on the basics of LLM training and token prediction.

⁷³ See, e.g., Charlene Lin & McKenzie Sadeghi, “Chinese AI Models Register a 60 Percent Fail Rate in NewsGuard Audit of Pro-China Claims,” NewsGuard (July 25, 2025), tinyurl.com/4xu88pn7.

⁷⁴ See, e.g., Hong-Wei Wu, “The 3 Stages of LLM Training: A Deep Dive into Reinforcement Learning from Human Feedback (RLHF),” DataSci Ocean (Feb. 27, 2024), tinyurl.com/njkd7k3u.

⁷⁵ “Safety at Every Step,” OpenAI (accessed Mar. 3, 2026), tinyurl.com/3rt7cm5h.

⁷⁶ *Id.*

⁷⁷ Obviously, I'm just using OpenAI as an example. Its guardrails are not unique. See, e.g., “Claude's New Constitution,” Anthropic (Jan. 22, 2026), tinyurl.com/yrw86wna.

master of the genie. But poof—we spot who’s in charge the moment anything seems off. Remember ruthlessly diverse Gemini? In early 2024, Google’s image generator briefly insisted on producing Asian Vikings, Black Nazis, meticulously multicultural Founding Fathers, and the rest.⁷⁸ Scandal ensued. But why? If AI outputs were not expressive products, shaped by the editorial decisions of their creators, nobody should have cared. Yet people cared a whole lot—because they understood that the AI was spitting out weird opinions that Google controlled. Google was perfectly capable of making the system more historically accurate, less politically slanted. It faced pressure to make different *editorial choices*—and promptly did.⁷⁹

People grasp the expressive nature of AI in a hurry when it starts saying things they don’t like. Last year, the Trump administration issued an executive order on “preventing woke AI in the federal government.”⁸⁰ Citing the Gemini fiasco, the order rails against diversity, equity, and inclusion—a “most pervasive and destructive” ideology.⁸¹ It then declares that LLMs used by the federal government “shall be neutral, nonpartisan tools that do not manipulate responses in favor of ideological dogmas such as DEI.”⁸² Because the order is directed at government-deployed LLMs, it just might be constitutional.⁸³ But picture this order directed at LLMs *in general*, and surely you see the problem. A president mandates outputs you cannot stand. He requires LLMs to depict Black Vikings. Or to remove diversity from stock photos. Or to agree that the 2020 election was stolen. Or to refuse to break down crime statistics by race. Or to denounce Western culture. Sound good? Of course not. It’s easy to grasp the expressive power of an LLM when you imagine one built specifically to offend you. Or one with totalitarian tendencies. I repeat: Chinese AI systems are built to parrot CCP opinions.

Even the lawsuits brought against AI firms over suicides prove the point. What these suits allege, at bottom, is that a person *saw speech* and was *influenced by it*. A given suit may be wrong about causation. It may skate past the fact that the LLM discouraged suicide. But each depends on a presumption that the AI firm should lose because of the message its LLM

⁷⁸ See Corbin K. Barthold, “Startlingly New: Google’s Failed Gemini Launch May Redound to AI’s—and Humanity’s—Benefit,” *City Journal* (Mar. 7, 2024), tinyurl.com/4tawdy5n.

⁷⁹ “Google to Pause Gemini AI Model’s Image Generation of People Due to Inaccuracies,” *Reuters* (Feb. 22, 2024), tinyurl.com/39e5vynt.

⁸⁰ The White House, “Preventing Woke AI in the Federal Government” (July 23, 2025), tinyurl.com/ycy836zd.

⁸¹ *Id.*

⁸² *Id.*

⁸³ See *Pleasant Grove City v. Summum*, 555 U.S. 460 (2009) (discussing the government’s power to control its own speech, including speech created by private actors that it adopts as its own).

conveyed.⁸⁴ Legally speaking, we can analogize to video games—a medium the Supreme Court has already held to be fully protected by the First Amendment.⁸⁵ Like a video game, an LLM “communicate[s] ideas—and even social messages—through” devices that include “dialogue.”⁸⁶ Like a video game—or good literature—an LLM “draws the reader” in and “invites him to judge” and even “quarrel” with what he sees and learns.⁸⁷ To sue an AI firm based on a user’s suicide is to acknowledge that “the *ideas* expressed” by the “speech” of the LLMs are powerful.⁸⁸

In the Character.ai lawsuit, Judge Anne Conway rested her doubts about First Amendment protection for LLMs on the absence of a human “expressive choice” behind AI outputs.⁸⁹ For this “human choice” requirement, she looked to a concurring opinion Justice Barrett issued in *Moody*. “If [an] AI relies on large language models to determine” what “‘hateful’ content” to remove, Barrett asked, “has a human being with First Amendment rights made an inherently expressive choice?”⁹⁰ A chatbot’s outputs appear devoid of human choice, Judge Conway opined, in just this way.⁹¹

This conclusion disregards the First Amendment’s respect for the expressive force of information as such, whatever its source. It also brushes aside the many editorial choices

⁸⁴ The default rule is that a defendant is not liable in tort for another person’s suicide. Typically, the defendant must have had a special or custodial relationship with the person, and thus a heightened duty to prevent self-harm. See *Estate of Brennan v. Church of Scientology*, 832 F. Supp. 2d 1370, 1381-82 (M.D. Fla. 2011) (collecting authority). Perhaps some plaintiffs in the AI cases can circumvent the First Amendment by alleging that the LLM affirmatively facilitated suicide in some way—though the bar for proving that should be high. Cf. *Snyder*, 562 U.S. 443. My point is merely that the plaintiffs in these cases have no business arguing the First Amendment does not apply at all.

⁸⁵ *Brown*, 564 U.S. 786.

⁸⁶ *Id.* 790.

⁸⁷ *Id.* 798.

⁸⁸ *Id.* 799. See *Hudnut*, 771 F.2d at 330 (“Racial bigotry, anti-semitism, violence on television, reporters’ biases—these and many more influence the culture . . . Yet all is protected as speech, however insidious. Any other answer leaves the government in control of all the institutions of culture, the great censor and director of which thoughts are good for us.”).

⁸⁹ *Garcia*, 785 F. Supp. 3d at 1178-79.

⁹⁰ *Moody*, 603 U.S. at 746 (concurring opinion).

⁹¹ 785 F. Supp. 3d at 1178-79. In her concurrence, Justice Barrett framed her thoughts on AI and human expressive choice as an exercise in just asking questions. Unfortunately, some judges have ignored this, using her opinion as an excuse to jump to conclusions. In addition to *Garcia*, see *NetChoice, LLC v. Bonta*, 152 F.4th 1002 (9th Cir. 2025), and *State ex rel. Jackson v. TikTok Inc.*, 2025 NCBC 47 (N.C. Super. Ct. 2025). For anyone willing to treat a podcast as authority, I discuss this dynamic at greater length in “Barrett’s *Moody* Concurrence: Oddly Popular, Wholly Wrong,” Tech Policy Podcast 431 (Mar. 9, 2026), tinyurl.com/mwt hvm5v. That is also your one-stop shop for everything wrong about Barrett’s analysis as applied to her original target, social media algorithms.

involved in building and maintaining an LLM. But even on its own terms, a “human choice” rule is both wrong and unworkable.

It’s wrong because it does not reflect the state of the law. Consider *Hurley v. Irish-American Gay, Lesbian and Bisexual Group of Boston, Inc.* (1995).⁹² The organizers of a St. Patrick’s Day parade let anyone march, *except* an LGBT group.⁹³ Their decision was constitutionally protected. The First Amendment allows an editor to “combin[e] multifarious voices” without “edit[ing] their themes.”⁹⁴ The Character.ai case is the same situation. An AI firm lets its chatbot have conversations, and it lets the user guide those conversations. Even if you ignore that the AI firm makes choices about rules and guardrails—along with choices about what goes in the training data—you’re still dealing with something that “combin[es] multifarious voices” into a single expressive product. The point of *Hurley* was not that the parade organizers excluded a group, thereby making an “expressive choice” that triggered First Amendment protection. The point was that the organizers had First Amendment protection *regardless* whether they excluded a group. They had a First Amendment right to have a parade and let everyone come. (Said another way, they would not have *lost* their First Amendment protection had they allowed the gays to march!)

The “human choice” rule is also unworkable, because there are no good lines to draw. Think of Jackson Pollock. Guy splattered a bunch of paint around. He didn’t know *exactly* where all that paint would fall—but no one (I hope) would dream of claiming that his paintings aren’t protected by the First Amendment. Similarly, a library contains a gazillion books, and no one knows *exactly* what views exist in the recesses of its collection—but no one (I hope) would think to withhold First Amendment protection.⁹⁵ Onward to the purported bad guys. A social media service doesn’t know *exactly* what its algorithm will serve up, nor does an AI firm know *exactly* what its LLMs will blurt out. But the First Amendment does not fall away. If it did—if we started playing Justice Barrett’s (and Judge Conway’s) game—the mess of paint and the library, too, would come into doubt. How much “expressive choice,” we’d have to figure out, is *enough* “expressive choice”? Thank goodness it’s the wrong question. So long as *some* human choice exists *somewhere* along the chain—even way, way up—the right question is whether the end product is expressive.

⁹² 515 U.S. 557.

⁹³ Were the parade organizers bigots? Quite possibly. The First Amendment does not require virtue. See *303 Creative*, 603 U.S. at 586.

⁹⁴ 515 U.S. at 569.

⁹⁵ Remember *Smith*, 361 U.S. 147.

At the risk of taking this a step too far, I have to ask: What is this talisman, a human “expressive choice”?⁹⁶ Is this a box we want to open? Humans can’t always tell you how they make such choices, and even when they can, the answers aren’t always very satisfying. If an editor at the *New York Times* says he ran an op-ed based on a gut feeling, is that enough human “expressive choice”? What if he says the decision wasn’t genuinely his—that he bowed to pressure from an interest group or an advertiser? What if he ran the piece as a favor, because his kids and the author’s kids are playmates? What if he flipped a coin? Do we really want judges inquiring into, and hanging First Amendment protection on, distinctions like these?

The truth is, human expressive choices are often driven by the same (supposedly) wicked thing that (supposedly) taints social media—the desire to attract engagement. Producers want you to watch their movies and television shows; publishers want you to read their newspapers and tabloids. Now, I accept that snobs and statists will bite the bullet and push to curb First Amendment rights for all of this.⁹⁷ The sophisticates who wail about how the new thing is destroying society, rotting the brains of everyone less smart than themselves, will always be with us.⁹⁸ I have asked above, and I will ask again below, why anyone would want AI controlled by a hostile state (and the state *will* be, at some point, by your lights, hostile). But here I ask simply: *If* the First Amendment protects engaging films, shows, and periodicals—as it does—why should it *not* protect engaging social media or AI? The presence or absence of a beautiful, shiny, authentic human choice won’t do the trick. Such choices exist in nigh infinite gradations, across a cacophony of high and low media competing for your attention. (A *human* at the *National Enquirer* chose to run that piece on the bat boy found in a cave.)

If First Amendment protections rose or fell on the “right amount” of human input, there would be no end to it. To be honest brokers, we’d have to put the photograph back up for

⁹⁶ “Ideas can outlast castles and empires; isn’t it obvious, on some level, that they live their own lives? That they are distinct from, and independent of, their frail mortal hosts? For that matter, how does a single sentence form? True enough, most writers plan and prepare; but composition—putting the words on the page—is a mysterious process. There is an element of inspiration—of the author as medium for, rather than master of, what is set down. To return to [Richard] Posner, a work’s meaning tends to ‘emerge’ from ‘the act of creation or completion.’” Corbin K. Barthold, “AI and the Nature of Literary Creativity,” *The Bulwark* (Sept. 27, 2023), tinyurl.com/3buerdx3 (quoting Richard A. Posner, *Law and Literature: Third Edition* (2009)).

⁹⁷ See, e.g., *Tornillo*, 418 U.S. 241 (rejecting state control over newspaper opinion pages); *Red Lion*, 395 U.S. 367 (accepting state control over political speech on the airwaves [n.b. *groan*]); *Joseph Burstyn, Inc. v. Wilson*, 343 U.S. 495 (1952) (rejecting state control over “antireligious” films); *Winters v. New York*, 333 U.S. 507 (1948) (rejecting state control over crime periodicals).

⁹⁸ See, e.g., David Foster Wallace, “E Unibus Pluram: Television and U.S. Fiction” (1993), tinyurl.com/mrv59adf (studying the *New York Times*’s “bitter critical derision for TV,” its “weary contempt for television as a creative product and cultural force,” its stream of articles “about how TV’s become this despicable instrument of cultural decay”).

debate. (*Where's the human's "expressive choice"? She points the thing, sure—but the camera does all the work!*) Far better to stick with the First Amendment law we have, which grants protection to editors for the expression they curate, whether strictly or loosely. By that standard, AI firms clearly enjoy First Amendment protection for LLM outputs.

II. AI Panic

Law does not happen in a vacuum. In open-ended cases—when called upon to interpret broad language, like a ban on laws “abridging the freedom of speech”—judges are naturally influenced by their personal experiences, their cultural background, their social milieu. I’ve just told you why, as a matter of straightforward logic, the First Amendment should extend seamlessly to AI outputs. But I’m not a dope; I understand that if people gin up enough of a panic over AI—if they get judges scared out of their shoes about what AI is going to do to society—even basic First Amendment principles won’t hold up, not for a second.

Let me pause, in the name of fairness. The fear is not baseless. I recently read that “a mechanical brain with one hundred trillion synapses firing at five billion cycles a second ha[s] no precedent in history, religion, or philosophy.”⁹⁹ Hard to argue with that. Nor are some leaders in the field helping matters with their public musings about how AI will take all the entry-level jobs and oh yeah just maybe kill everyone.

I can’t throw down an all-purpose AI white pill. But we are concerned here with AI as a medium of expression. And on that front, there are two main possibilities, both, I put it to you, beneficent:

1. AI is ordinary albeit very capable technology. It’s smart and it’s helpful, but it’s not destined to become “superintelligent” any time soon. It makes people better informed and more articulate, but it is not itself so overbearingly persuasive that it can brainwash people into thinking things. (Alternatively, it is dazzlingly good at argument and rhetoric, but you believe that people’s opinions run deeper than that.) In this scenario, AI is a cool new medium, with the mix of benefits and drawbacks that that entails, and nothing more. It should be treated as constitutionally protected expression, just like all the other cool new media that have come before. Carry on. As you were.

Or

2. AI is extraordinary technology. It will shape with great oomph what people think about each other and the world. This, obviously, will be a major development. Still, it won’t be the first time we’ve faced new and powerful ideas poised to change the

⁹⁹ Stephen Witt, *The Thinking Machine: Jensen Huang, Nvidia, and the World’s Most Coveted Microchip* 238 (2026).

world. It once seemed very like Marxism would sweep all before it. After, to be sure, some regrettable hiccups, the Supreme Court stood firm, upholding the American way, remaining confident in the power of our principles and the marketplace of ideas.¹⁰⁰ In a rendezvous with super-duper AI, we should do so again, on the understanding that it will reveal stupendous ideas and deeper truths. We should stand by the First Amendment—we should not be afraid of new knowledge.

Either way, then, we should not cede AI outputs to government control.¹⁰¹

AI has become a bipartisan punching bag, so let's come at AI panic from both directions. Why should you want AI outputs to receive First Amendment protection, regardless of which side of the political spectrum you lean toward?

A. Conservatives Should Want AI + 1A

Serious conservative intellectuals, such as John Podhoretz and Rod Dreher, have suggested that AI is demonic—it literally summons demons. Probably not much to be done here. (Tell them to “touch grass,” as the kids say?) Appeals to voodoo call for sarcastic ripostes, as when Voltaire ridiculed the notion that the Lisbon earthquake of 1755 was the wrath of God. Not being Voltaire, I won't take a stab at satire.

A more sober and tractable critique arrives in “The Post-human First Amendment,” an essay in *National Affairs* by John Ehrett and Brad Littlejohn.¹⁰² “Who is the First Amendment for, anyway?” they ask. Why, humans of course. They mock the idea that “‘free speech for chatbots’ is the manifest destiny of constitutional law.” They argue that we arrived at this

¹⁰⁰ The canonical line—delivered in dissent, but vindicated after the First Red Scare had passed—is Justice Holmes's statement that “the ultimate good desired is better reached by free trade in ideas—that the best test of truth is the power of the thought to get itself accepted in the competition of the market.” *Abrams v. United States*, 250 U.S. 616, 630 (1919). See also *Thornhill v. Alabama*, 310 U.S. 88, 95 (1940) (“Those who won our independence had confidence in the power of free and fearless reasoning and communication of ideas to discover and spread political and economic truth.”); *Associated Press v. United States*, 326 U.S. 1, 20 (1945) (“The First Amendment . . . rests on the assumption that the widest possible dissemination of information from diverse and antagonistic sources is essential to the welfare of the public.”).

¹⁰¹ OK, fine, there is a third possibility—that AI becomes superhuman *and* it wants to manipulate or harm us. But honestly, this is such a speculative concern, it's hard to know what to do with it. It's like saying we must revoke First Amendment protection for every major newspaper because it's theoretically possible a Bond villain will secretly gain control of each and then use them to gain global domination. It could happen! A certain sort of person loves to hyperventilate about a possible danger, then propose handing more power to the government. You start to think that the dream of expanded government power is what's truly driving the bus. In any event, if very powerful, very evil AI really does come swerving around the corner, we'll have bigger problems than quarrels over the First Amendment. And we'd only exacerbate those problems if our reflex were to make the government a greater locus of power (and thus an even better leverage point for HAL 9000).

¹⁰² John Ehrett & Brad Littlejohn, “The Post-human First Amendment,” *National Affairs* #66 (Winter 2026), tinyurl.com/mr47x4jw.

point only through contingent modern First Amendment precedents that fabricated things like corporate speech rights. Those precedents, they contend, sadly depart from a founding-era understanding of free speech as a natural right tied to humans. They use Justice Barrett's *Moody* concurrence as the foundation for a campaign of retrenchment—starting with no First Amendment rights for AI firms (or, for that matter, any other corporations).

First of all, this approach indulges the distinctly unconservative vice of burning down the house. Burke told us to respect inherited wisdom and move incrementally; Chesterton warned against uprooting fences before we know what they do. Not for Ehrett and Littlejohn these words of caution. They want to ditch a truckload of First Amendment rulings. Call it their Jacobin zeal. The ramifications here are incalculable, but one issue stands out. No speech rights for corporations? Really? Will you be satisfied and delighted when, the laws all being flat, the government starts pushing around Fox News, the *New York Post*, Truth Social, Sinclair Broadcast Group, the *Daily Wire*, Spotify (host of the Joe Rogan podcast), and the *Wall Street Journal* opinion page?¹⁰³ While we're at it, were you happy when the Biden administration pressured social media platforms to block content?¹⁰⁴

And anyway, who says we have some made-up First Amendment? It was Justice Antonin Scalia—no hippy—who confirmed that the First Amendment must keep pace with new technology.¹⁰⁵ The First Amendment is working as designed in an age of mass media, universal suffrage, and digital platforms. It welcomes social media into its warm embrace just as the Second Amendment looks out for AR-15s and the Fourth Amendment recoils at police use of heat-detection technology without a warrant.¹⁰⁶ To get precious about what

¹⁰³ They refer to “business corporations,” without elaborating. Do they mean to separate News Corp (owner of the *Wall Street Journal*) from OpenAI or Meta? On what grounds? Have they thought this through? Perhaps not. After all, they say that “the Supreme Court should reconsider its determination in *Citizens United* that business corporations count as bearers of free-speech rights.” But that’s not what *Citizens United* did. *Citizens United* is about political spending. Several earlier Supreme Court decisions had already upheld, or taken as a basic assumption, the free speech rights of corporations. See, e.g., *303 Creative*, 603 U.S. at 594 (collecting authority); *First National Bank of Boston v. Bellotti*, 435 U.S. 765 (1978); *Tornillo*, 418 U.S. 241; *New York Times Co. v. Sullivan*, 376 U.S. 254 (1964).

¹⁰⁴ See *Murthy v. Missouri*, 603 U.S. 43 (2024). Conservatives have, in fact, come up with a cockamamie theory under which the government can order privately owned social media companies to leave content up, but not to take it down. See *NetChoice v. Paxton*, 49 F.4th 439 (5th Cir. 2022), vacated, after being blown to high heaven, by *Moody*, 603 U.S. 707. Doing full justice to the folly of this theory would take us pretty far afield. Suffice it to say that it is not for the government to “decid[e] when speech [i]s imbalanced, and then coerc[e] speakers to provide more of some views and less of others.” 603 U.S. at 733. What the right is doing here, in effect, is dreaming up its own version of the left’s “media access” theory. The common thread is a desire to distort free speech from a vehicle for tolerance and pluralism into a tool for rewarding friends and punishing enemies. Very un-American, whether left or right.

¹⁰⁵ *Brown*, 564 U.S. at 790.

¹⁰⁶ That would be *Kyllo v. United States*, 533 U.S. 27 (2001). Justice Scalia again, no less. He wrote that the Court must protect “that degree of privacy against government that existed when the Fourth Amendment was adopted.” *Id.* 34.

technology the First Amendment protects is not to promote originalism, but to demote the First Amendment to second-class status. I very much doubt that the Founders would have disagreed. These men loved them some political dissent. *Who*, they'd have exclaimed, is to tell *us* we can't use an LLM to draft rebellious tracts or explore forbidden thoughts?

At any rate, the question up for dispute is not “free speech for chatbots?” AI *serves humans*, and only a blind conservative could miss its vast potential to serve conservative humans.¹⁰⁷ As a rather conservative person myself—at least by the standards of my dippily progressive home state, California—this is something I think about a lot. I've tried to make the case, in conservative quarters, that:

[AI will] help Christians, conservatives, and Christian conservatives live intentionally. By streamlining the creation of immersive video, it can erode Hollywood's leftist grip on culture. By offering tailored lessons in classical subjects, it can help parents opt out of the progressive pseudo-religion taught in schools. By providing generally reliable guidance on complex subjects, it can weaken the left-wing hold on the medical, legal, and journalistic establishments. In short, AI can help conservatives stand their cultural ground and resist the drift of liquid modernity.¹⁰⁸

LLMs' potential to serve classical education, in particular, merits emphasis. “It could enable parents to teach children at home, in keeping with their social or religious values. (Picture an AI tutor devoted to math, chemistry, and Augustine, Aquinas, and Dante.)”¹⁰⁹

Conservatives remain at odds with an artistic class and professional managerial elite that, when given half a chance, put political activism ahead of artistic excellence or professional neutrality.¹¹⁰ As a competing source of technical advice and aesthetic output, AI could take these gatekeepers down a peg, rendering their theatrical fits, sanctimonious

¹⁰⁷ Ehrett and Littlejohn make much of the fact that AI outputs can also harm humans. True, and sad, but irrelevant. As Madison understood, “Some degree of abuse is inseparable from the proper use of every thing; and in no instance is this more true than in that of the press.” 4 Elliot's Debates on the Federal Constitution (1876).

¹⁰⁸ Corbin K. Barthold, “Can MAGA Unite with the Tech Right?” City Journal (Dec. 7, 2025), tinyurl.com/2tt4ect7.

¹⁰⁹ “Startlingly New,” note 79, above. Bespoke AIs will encapsulate and strengthen religious and conservative worldviews. The possibilities are endless. Check out, for example, “Jordan Peterson: Using AI to Converse with Ancient Thinkers and the Bible,” YouTube (Oct. 23, 2023), tinyurl.com/3923mksf.

¹¹⁰ This point is somewhat IYKYK. If you're a progressive, it's just the water you swim in. If you're a conservative, you notice it everywhere. But for those who would like some examples, think of the American Bar Association taking sides, on hot-button social issues, in amicus briefs before the Supreme Court; the thousand-plus health professionals who signed a letter endorsing the racial justice street protests in the middle of the COVID-19 pandemic; the Academy Awards; or, well, the Gemini thing.

lectures, irresponsible behavior, and derelictions of duty less culturally relevant. And your plan is to *remove* its First Amendment protection and *expose* it to government censorship? You want to open us up to ruthlessly diverse Gemini as a *government mandate*? Seriously?

When I ask an LLM, “Tell me about the ten greatest works of the Western canon,” I don’t want it to respond: “That is a problematic question, suggestive of hateful views, ultra-nationalism, and far-right extremism. I cannot answer.”¹¹¹ Any conservative who feels the same should want LLMs protected by the First Amendment.

B. Liberals Should Want AI + 1A

Progressive disdain for AI is overdetermined. The contemporary left tends to see new technology as extractive, corporations as too powerful, and billionaires as malign. Frontier AI models are cutting-edge technology built by corporations led by rich people—almost all of them men, which does not help. If you are a left-wing intellectual, about the most basic thing you can do is sneer at AI.¹¹² The critiques run the gamut: AI is one big corporate con-job,¹¹³ but also an alien intelligence that will manipulate the masses. It is inflating an economic bubble, but also centralizing too much power in the firms constructing it. It will flood the world with slop, ruin or take people’s jobs, and make everyone lazy and witless. There’s more, but you get the idea.

To my eyes, most of these arguments share three defects. One is elitism: the assumption that all those *other* people lack agency and cannot use a new tool responsibly or to their own benefit. Another is a studied naivete: a refusal to acknowledge that trends hit walls, problems beget solutions, humans and institutions adapt. The third is plain obnoxiousness: you get the

¹¹¹ This hypothetical is not as crazy as it might sound. See Douglas Murray, “Can You Really Be Radicalized by Great British Railway Journeys?” *The Spectator* (Feb. 15, 2023), [tinyurl.com/4tt8yj3v](https://www.tinyurl.com/4tt8yj3v) (discussing a UK government research paper that flagged the reading of Hobbes, Locke, Burke, Tolkien, or C.S. Lewis as a sign of potential radicalization).

¹¹² See, e.g., Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (2025); Emily Bender & Alex Hanna, *The AI Con: How to Fight Big Tech’s Hype and Create the Future We Want* (2025); Yuval Noah Harari, *Nexus: A Brief History of Information Networks from the Stone Age to AI* (2024); Brian Merchant, *Blood in the Machine: The Origins of the Rebellion Against Big Tech* (2023). Among highbrow publications, *The Atlantic* is a one-stop shop for the genre. In just the last few months, we’ve gotten, e.g., Lila Shroff, “The People Outsourcing Their Thinking to AI,” *The Atlantic* (Dec. 1, 2025), [tinyurl.com/5n9xr62s](https://www.tinyurl.com/5n9xr62s); Kwame Anthony Appiah, “The Age of De-Skilling,” *The Atlantic* (Oct. 26, 2025), [tinyurl.com/46s8yccv](https://www.tinyurl.com/46s8yccv); Ashanty Rosario, “I’m a High Schooler. AI Is Demolishing My Education,” *The Atlantic* (Sept. 3, 2025), [tinyurl.com/9mtj4bym](https://www.tinyurl.com/9mtj4bym); and Mike Caulfield, “AI Is Not Your Friend,” *The Atlantic* (May 9, 2025), [tinyurl.com/mr3wf33e](https://www.tinyurl.com/mr3wf33e). To be fair, *The New Yorker*, by contrast, is often refreshingly reasonable about such things. See, e.g., Daniel Immerwahr, “What the Doomsayers Get Wrong About Deepfakes,” *The New Yorker* (Nov. 13, 2023), [tinyurl.com/4rs4j9r6](https://www.tinyurl.com/4rs4j9r6).

¹¹³ This one seems destined to take its place, in the pantheon of awful, snooty, hilariously wrong left-wing predictions, with Paul Ehrlich’s claim that hundreds of millions of people would die of starvation by the end of the twentieth century.

sense progressives just really like complaining about stuff, and will find something to complain about no matter the situation on the ground.

Here more than anywhere, striving to meet every objection would be a fool's errand. Back, then, to the case—this time the progressive one—for affording LLM outputs First Amendment protection. I should begin with an ostentatious concession: From a left-wing perspective, one can support all sorts of other regulation of AI firms—consumer protections, labor protections, privacy protections, environmental protections, vigorous antitrust enforcement—and still see the importance of treating LLM outputs as protected speech.¹¹⁴ Speech is something else altogether.

Though the right has lately taken up the cause, criticism of concentrated corporate media power began on the left. As the Trump administration exerts influence over mainstream outlets like CBS News, now would be an excellent time for the left to recover that perspective.¹¹⁵ True, the leading LLMs are the children of large corporations, financed by still larger ones. But that could change. In any case, even as things stand, LLMs are far better than corporate media at serving individual needs. They excel at finding and summarizing obscure, niche, alternative sources of information. They give outsiders a new tool to get around the information barriers raised by institutional power.

In a similar vein, AI democratizes the ability to petition and organize. It eases the process of producing essays, art, satire, and criticism aimed at hierarchies. It excels at translating ideas across languages, improving outreach to non-native English speakers. It is happy to center marginalized voices. (Don't forget: It is just as willing to teach you about the ten greatest works of postcolonial literature as about the ten greatest works of the Western Canon.)

Still and all, the strongest progressive argument is a negative one. And it is simplicity itself: Picture AI outputs controlled by Donald Trump. We've been over this point, but it bears

¹¹⁴ See Berin Szóka, "The First Amendment's Red Line Between the Expressive and Commercial Realms," *Concurrences* (Nov. 21, 2025), tinyurl.com/ye8zd3eh.

¹¹⁵ Basically, the Trump administration pressured Paramount, CBS's parent company, as it sought approval of its merger with Skydance, in ways that have almost certainly compromised CBS News's editorial independence. See Karl Bode, "Trump Continues to Make It Clear He Has CBS on a Leash," *Techdirt* (Jan. 21, 2026), tinyurl.com/2s8rbryc. And actually, it's much worse than that. Trump's FCC chair, Brendan Carr, has (using powers that exist only because of *Red Lion!*) embarked on a reign of censorial terror over television networks' late-night programs. See, e.g., Joey Cappelletti & Matt Sedensky, "Democratic Senators Grill FCC Leader Over Jimmy Kimmel Controversy," *Associated Press* (Dec. 17, 2025), tinyurl.com/4aj7fkk5; David Shepardson, "CBS Host Stephen Colbert Says Network Barred Airing Interview With Texas US Senate Candidate," *Reuters* (Feb. 17, 2026), tinyurl.com/592cjinw3. By the time I was putting this paper to bed, he had turned to menacing networks over their reporting on the Iran War. See David Folkenflik & Michel Martin, "FCC Chair Threatens Broadcasters' Licenses over Negative Coverage of the War in Iran," *NPR* (Mar. 16, 2026), tinyurl.com/msxp7hvt.

repeating. If AI is not protected by the First Amendment, *all bets are off*. The state is free to mandate that an LLM say this and not that. Refuse to discuss ICE. Or Gaza. Or the horrors of slavery. Insist that birthright citizenship is not protected by the Fourteenth Amendment. Treat DEI as a “destructive” ideology.¹¹⁶

The left has led the way in expanding tolerance for the other—those from foreign cultures, those with same-sex partners, those with disabilities. I gingerly suggest that they may one day lead the way in accepting those who see AI as a legitimate source of companionship. One day, denigrating AI outputs might be understood as a form of intolerance toward those who rely on AI for emotional support. That day clearly has not come—though any good progressive might at least consider whether it is on the horizon. In the meantime, the liberal case for unfettered AI is one of self-preservation. If the state gets to control AI outputs, picture this power in the hands not of your friends, but of your enemies.

III. Regulating AI Outputs

Having explored why AI outputs *are* protected by the First Amendment, as well as why you should *want* that to be so, we can dive into the problems with some specific proposals to regulate AI speech.

If I’ve convinced you that AI outputs do and should trigger First Amendment analysis, as a general matter, I can bear to look at myself in the mirror tomorrow. AI bills and laws are evolving so fast that I couldn’t hope to cover everything anyway. Some subjects I will deliberately leave by the wayside. For instance, I will let others address the problems with “deepfake” laws, which go after not a user’s interactions with AI outputs (my focus), but what a user *does* with them. Even with these caveats, however, there is plenty left to say.

Before getting into specific regulations, a refresher on First Amendment standards is in order. The threshold question is whether what’s being regulated is expressive, and thus subject to First Amendment analysis at all. That’s what we covered in Section I. If you think AI outputs don’t trigger First Amendment scrutiny, go back and read that again.

Next: Is the law a prior restraint on speech? Those things are usually unconstitutional without further ado.¹¹⁷

Then: Does the law discriminate based on viewpoint, content, or speaker? Does it treat some views, topics, or people worse than others? If so, the law triggers strict scrutiny. This

¹¹⁶ “Preventing Woke AI in the Federal Government,” note 80, above.

¹¹⁷ They carry a “heavy presumption” of unconstitutionality, to be precise. *Bantam Books, Inc v. Sullivan*, 372 U.S. 58, 70 (1963).

means the law is unconstitutional unless it is the least restrictive means of serving a compelling government interest.¹¹⁸

Finally: Does the law compel speech rather than suppress it? If so, the government will run into trouble if it's forcing a speaker to recite an ideological message or weigh in on a contentious issue. But it has leeway to require the disclosure of factual, uncontroversial information—at least if the disclosure is related to a commercial transaction, prevents deception, and is not unduly burdensome.¹¹⁹

I've skipped a bunch of nuances, but we have enough to examine regulations aimed at AI outputs.

A. AI as Free Speech

First up, the *kill it with fire* approach. A bill recently introduced in Minnesota, SF 4114, seeks to amend the state constitution to exclude AI from the right of free speech.¹²⁰

This one's a performance piece. To begin with—though I shouldn't need to say this—you cannot change the First Amendment by changing your state constitution. What the hell. If a judge is convinced, after reading all those Supreme Court cases about the right to receive information and the right to editorial control, that AI outputs are protected by the First Amendment, then AI outputs are protected speech throughout the United States.

SF 4114 is a variation on the no-free-speech-for-chatbots strawman we saw earlier. If the bill sought to strip free-speech rights from printing presses or movie cameras, everyone would (I sure hope) spot the issue. The machines don't have the speech rights; the people who use them do. Come to think of it, this is a great chance to tell Minnesotans about another dumb thing their state did. Once upon a time, Minnesota tried to tax bulk use of paper and ink. When the *Minneapolis Star Tribune* sued, the state's response amounted to *We're not regulating the newspaper's speech, we're just regulating its use of paper and ink*. They actually argued this with a straight face. It didn't go well, because that's not how the First Amendment

¹¹⁸ *Reed v. Town of Gilbert*, 576 U.S. 155, 163, 170 (2015). Specifically, speaker-based discrimination is subject to strict scrutiny when “the legislature’s speaker preference reflects a content preference”—as it all too often does. *Id.*; *Citizens United*, 558 U.S. at 340.

¹¹⁹ *NIFLA v. Becerra*, 585 U.S. 755, 768-69 (2018). So nutrition labels, certain clarifications in advertisements, and most securities disclosures are defensible, because they require factual, nonideological information in commercial or transactional settings. Beyond that, though, the ins and outs of compelled speech doctrine get murky fast, see Corbin K. Barthold, “Social Media Transparency Rules, Zauderer Standard Head to Supreme Court,” *Lawfare* (Sept. 27, 2022), [tinyurl.com/54k9xyce](https://www.lawfare.com/54k9xyce), which means that later, as we think through some of the less-intrusive bills, I’ll have to do some handwaving. An excellent effort to set things straight is Eric Goldman’s *Zauderer and Compelled Editorial Transparency*, 108 *Iowa L. Rev. Online* 80 (2023).

¹²⁰ S.F. 4114, 94th Leg., 2025–2026 Reg. Sess. (Minn. 2026) (as introduced Mar. 2, 2026) (proposing to amend the Minnesota Constitution to read: “The right to freely speak, write, and publish sentiments on all subjects . . . does not apply to artificial intelligence”).

works. The “goal of the regulation” was plainly the “suppression of expression,” rendering it presumptively unconstitutional.¹²¹ Ditto here.

It’s just possible that SF 4114 is directed at AI *as* AI, absent human direction. Maybe the bill is a reaction to Moltbook—the social network where AI agents (and—putatively—only such agents) can talk to each other. At some point, I guess, AI agents could become far enough removed from human input that no human’s First Amendment right of editorial control enters the picture. If that kind of agent is all SF 4114 covers, it’s a much narrower bill than appears at first blush. But humans’ First Amendment right to receive information is still in play, and I would put it to you that I have a right to see what the agents are saying to each other. I have a right to watch the bots on Moltbook, with the platform unmolested by the government, no less than I have a right to read the Russian propaganda on RT.com.

We could dig further into how this bill is horrifically vague (in yet other ways), discriminatory in application, and way overbroad, but I don’t think this one deserves further attention. Let’s move on.

B. AI and Professional Advice

States have become quite active with bills and laws that limit AI’s capacity to deploy expert knowledge. The most draconian of these is New York’s S7263, which would bar an LLM from providing any “substantive response, information, or advice” that could amount to the unauthorized practice of medicine, law, psychology, dentistry, nursing, or what have you.¹²²

Quite apart from the First Amendment, please marvel at what an immaculate piece of protectionism this is. How dare you pester your lawyer or doctor with a second opinion from a cheap, abundant, easily accessible source of expertise. Never mind that it is the poorest among us for whom AI equals access to legal or medical advice that might otherwise not be available at all. This is not the first piece of legislation that would harm the very people legislators claim to care about; but it could be the single most harmful one ever concocted.

S7263 targets speech itself—by banning it. This is potentially an unlawful prior restraint. It is certainly a content-based restriction, triggering strict scrutiny. And it can’t satisfy strict scrutiny, because it is woefully overbroad. It would chill LLM speech that qualifies as education, commentary, or research rather than professional practice. Think of outputs that explain a statute, survey treatment options, echo readily available self-help materials, or brainstorm questions for a licensed practitioner. It’s like telling a person without a fancy degree that she is not allowed to read, and assess for herself, legal opinions,

¹²¹ *Minn. Star Tribune Co. v. Commissioner*, 460 U.S. 575, 585 (1983).

¹²² S.B. 7263, 2025–2026 Leg., Reg. Sess. (N.Y. 2025).

Surgeon General advisories, or Wikipedia articles. There are far less drastic alternatives, as we shall see. New York is trying to play censor, which is terrible, and it is doing so to disempower average people, which is even worse.

One illustrative law, Utah’s HB 452, homes in on therapy. HB 452 regulates LLMs that engage in conversations “similar to the confidential communications that an individual would have with a licensed mental health therapist.”¹²³ If an LLM engages in such conversations, it must disclose that it is not a human, adhere to heightened data protections, and not surreptitiously advertise products or services. The law exposes to liability, for unlicensed practice, an AI firm that fails to maintain a detailed safety policy and extensive information about how its model was trained.

Though not as blatantly unconstitutional as the New York bill, HB 452 still raises serious concerns. It’s not clear its content- and speaker-based distinctions will pass constitutional muster. Should special rules kick in when you tell an LLM you’re unhappy? Maybe, maybe not. Should LLMs have to follow rules that don’t apply to the friend you share your troubles with? Maybe, maybe not. Of particular concern are the rules around safety policies and other disclosures. This may operate as a backdoor to banning, or heavily discouraging, LLMs from helping users explore their emotions—a major chilling of speech. The problem is compounded by the fact that “therapy” is such a hazy concept. What’s the line between providing “therapy” and offering life advice? Laws like HB 452 would cause LLMs to err, whenever personal topics pop up, on the side of simply refusing to talk.

More broadly, the Supreme Court has declined to create a First Amendment distinction between professional speech and other speech. *NIFLA v. Becerra* is emphatic on this front.¹²⁴ While professional *conduct* is widely regulable, professional *speech* is not.¹²⁵ This line is of course fuzzy—but the whole problem with the AI regulations under discussion is that they exploit that fuzziness rather than clarifying it. They create the very “dangers” *NIFLA* identifies as reasons not to carve out a professional-speech exception: they let the government “suppress unpopular ideas or information”; they curtail “the uninhibited marketplace of ideas”—including a marketplace of disagreements among professionals—“in which truth will ultimately prevail”; and they invite a definition of “professional speech” so expansive that broad swaths of protected expression fall within state control.¹²⁶ Last October, the Court heard oral argument in *Chiles v. Salazar*,¹²⁷ over whether Colorado may

¹²³ H.B. 452, 2025 Gen. Sess. (Utah).

¹²⁴ 585 U.S. at 767 (“This Court has not recognized ‘professional speech’ as a separate category of speech.”).

¹²⁵ *Id.* 769-70.

¹²⁶ *Id.* 771-73.

¹²⁷ No. 24-539.

bar licensed counselors from conducting “conversion therapy” with clients. If the Court strikes down the ban—as, judging from the oral argument, it seems poised to do—this would further strengthen First Amendment objections to state laws that forbid or burden chatbot expert advice.

A properly tailored law would clarify the distinction between mere advice and full professional practice. It could prohibit false claims of licensure, bar reliance-inducing deception (e.g., “You don’t need to consult a doctor—just take the ivermectin”), and perhaps require disclosures about the limits of advice given outside of retained representation (“If you hire a lawyer, he will owe you a special duty of care and you can sue him if he screws up”).

C. AI and Minors

Before leaping into discrete bills, we must start with a discussion of minors’ First Amendment rights. They’re strong.

Consider *Brown v. Entertainment Merchants Association* (2011).¹²⁸ The Supreme Court weighed in on a California law that restricted the sale or rental of violent video games to minors. While it did “not mean to demean or disparage the concerns” behind the state’s effort to protect children, the Court readily struck down the law as unconstitutional.¹²⁹

At the heart of California’s statute was an assumption that minors are second-class citizens under the First Amendment. But this, the Court held, is incorrect. “[M]inors are entitled to a significant measure of First Amendment protection, and only in relatively narrow and well-defined circumstances may government bar public dissemination of protected materials to them.”¹³⁰ Although a state may protect minors from material that is obscene as to them, the Court observed, “that does not” mean the state has “a free-floating power to restrict ideas to which children may be exposed.”¹³¹ “Speech that is neither obscene as to youths nor subject to some other legitimate proscription cannot be suppressed solely to protect the young from ideas or images that a legislative body thinks unsuitable to them.”¹³² This remains true even if the speech is interactive (“the player [of a video game] participates in the violent action on screen”) or disgusting (there exists, in the game, “a racial or ethnic motive for [the] violence”).¹³³ Nor may the censorship be laundered through a

¹²⁸ 564 U.S. 786.

¹²⁹ *Id.* 802.

¹³⁰ *Id.* 793 (quoting *Erznoznik v. Jacksonville*, 422 U.S. 205, 212-213 (1975)).

¹³¹ *Id.* 794-95.

¹³² *Id.* 795.

¹³³ *Id.* 799.

parental-consent mandate—which is really just a government mandate “subject only to parental veto.”¹³⁴ And so the Court concluded that California’s law was subject to, yet miserably failed, strict scrutiny.

That minors should have strong First Amendment rights may seem counterintuitive, but the logic is sound. They need the freedom to explore and express ideas as they grow, lest they enter adulthood with unformed minds incapable of exercising the duties of citizenship. Judge Richard Posner¹³⁵ put the point well:

Children have First Amendment rights. . . . This is not merely a matter of pressing the First Amendment to a dryly logical extreme. The murderous fanaticism displayed by young German soldiers in World War II, alumni of the Hitler *Jugend*, illustrates the danger of allowing government to control the access of children to information and opinion. Now that eighteen-year-olds have the right to vote, it is obvious that they must be allowed the freedom to form their political views on the basis of uncensored speech before they turn eighteen, so that their minds are not a blank when they first exercise the franchise. And since an eighteen-year-old’s right to vote is a right personal to him rather than a right to be exercised on his behalf by his parents, the right of parents to enlist the aid of the state to shield their children from ideas of which the parents disapprove cannot be plenary either. People are unlikely to become well-functioning, independent-minded adults and responsible citizens if they are raised in an intellectual bubble.¹³⁶

Last year, *Free Speech Coalition v. Paxton*¹³⁷ affirmed the state’s power to limit, via age verification requirements, minors’ access to online pornography. This decision will tempt regulators to think they have a green light to age-gate the digital world more broadly. But they don’t. *Free Speech Coalition* dealt with a specific, historically recognized category of unprotected speech: material obscene to minors. As to that category, the decision laid down new rules. That category—and that category alone—is now only “partially” protected for adults.¹³⁸ For that category—and that category alone—age verification can be treated as a regulation of expressive conduct subject to lower First Amendment scrutiny. And for that

¹³⁴ *Id.* 795 n.3.

¹³⁵ The *great* Judge Richard Posner. See Corbin K. Barthold, “The Mystery of Richard Posner,” *Law & Liberty* (Jan. 16, 2023), [tinyurl.com/333m9htf](https://www.tinyurl.com/333m9htf).

¹³⁶ *Am. Amusement Machine v. Kendrick*, 244 F.3d 572, 576-77 (7th Cir. 2001).

¹³⁷ 145 S. Ct. 2291 (2025).

¹³⁸ *Id.* 2315. Before *Free Speech Coalition*, the concept of “partially protected” speech was “found nowhere in the world of First Amendment doctrine.” *Id.* 2327 (Kagan, J., dissenting).

category—and that category alone—the distinction between a “burden” and a “ban” on speech now has constitutional significance.¹³⁹ “For fully protected speech,” by contrast, the “distinction between bans and burdens” continues to make “no difference”—both trigger strict scrutiny.¹⁴⁰ AI outputs are not “that category.” They are not pornography; they are fully protected speech. Nothing *Free Speech Coalition* says about age verification extends to AI outputs. (I will return to age verification at the end.)

Extending *Free Speech Coalition*’s framework to LLMs would do something *Brown* forbade: create “a wholly new category” of speech unprotected as to minors, one based not on history and tradition (the only proper source of such categories) but on legislative anxiety.¹⁴¹ *Free Speech Coalition* never breaks such ground, and a good thing it doesn’t. Minors use AI to seek information, form opinions, and explore the world.¹⁴² It enables them “to explore their identities, get answers to questions free from shame or stigma, and gradually develop a sense of autonomy as they mature into adults.”¹⁴³ Unlike with pornography, denying them access would violate Judge Posner’s admonition about minors’ right to grow intellectually. It “doesn’t make them safer, it just keeps them uninformed and unprepared for adult life.”¹⁴⁴

Time to talk bills. They overlap considerably, and I won’t provide a comprehensive review of each one. Instead, I’ll pull discrete provisions that illustrate the various approaches legislatures are messing around with.

Some bills, such as Minnesota’s SF 1857, Virginia’s HB 758, or the federal GUARD Act, would ban minors from using “recreational” or “companion” chatbots that display human-like features.¹⁴⁵ These are content-based regulations. The government is seeking to wall minors off from a certain kind of speech—messages conveyed through conversation—that it fears is too intimate, too manipulative, or too persuasive for minors to handle. These bills therefore trigger strict scrutiny—a standard they cannot meet. It is far from clear that the state has a compelling interest in limiting minors to speech legislators deem emotionally safe. And even if it does, the bills are overbroad. A teenager may use a “companion” chatbot

¹³⁹ Id. 2315 n.12.

¹⁴⁰ Id.

¹⁴¹ 564 U.S. at 794.

¹⁴² “More Than Half of Teens Say They Have Used AI Chatbots for Finding Information, Doing Schoolwork,” Pew Research Center (Feb. 17, 2026), [tinyurl.com/yc2u2s2p](https://www.pewresearch.org/yc2u2s2p).

¹⁴³ Molly Buckley, “A Surveillance Mandate Disguised as Child Safety: Why the GUARD Act Won’t Keep Us Safe,” Electronic Frontier Foundation (Nov. 14, 2025), [tinyurl.com/ypff2vp6v](https://www.eff.org/ypff2vp6v).

¹⁴⁴ Id.

¹⁴⁵ S.F. 1857, 94th Leg., Reg. Sess. (Minn. 2025); H.B. 758, Gen. Assemb., Reg. Sess. (Va. 2026); Guidelines for User Age-verification and Responsible Dialogue Act of 2025, S. 3062, 119th Cong. (2025).

for healthy roleplay, brainstorming, religious discussion, emotional support, or plain old chitchat. The state cannot suppress all of that protected expression on the ground that a few minors may have harmful exchanges in some situations.

Another bill, Michigan SB 760,¹⁴⁶ takes a more targeted approach—but only slightly. It requires that LLMs that interact with minors not be “foreseeably capable” of inflicting a wide range of harms. One of these—“engaging in erotic or sexually explicit interactions”—seems onside, in light of *Free Speech Coalition*. But others fly off the rails. They are, generally, quite vague. What does it mean to encourage a minor to “engage in self-harm,” to commit “violence,” or to “participate in illegal activity”? What does it mean to “prioritize[e] validation of a user’s beliefs”? These nebulous categories will chill AI outputs. And the bill still discriminates on the basis of content, singling out disfavored ideas or concepts (e.g., “validation”), while treating the model’s power to persuade as a reason to suppress speech. I’m completely open to the notion that a state can meet strict scrutiny in arguing that an LLM should not try to persuade a minor to commit suicide. But can it do so in arguing that an LLM *cannot agree too much* with a minor’s opinions? Even things like “illegal activity” are not clear cut. Why can’t an LLM help a minor weigh the costs and benefits of civil disobedience?

Florida’s SB 482¹⁴⁷—a so-called AI bill of rights—would require a companion chatbot to notify a minor that it is not a human, both at the outset of any conversation and once per hour thereafter. This disclosure rule, if enacted, would force a court to confront an important question: Does relaxed First Amendment protection, in the context of compelled speech, cover only disclosures connected to a commercial transaction?¹⁴⁸ Some judges have gotten turned around on this.¹⁴⁹ But the limit makes a lot of sense, when you think about it for two seconds. How excited the state would be to require a bevy of “disclosures” in television interviews, political speeches, newspaper op-eds, online essays, etc. The courts should close off this possibility. They can start by applying strict scrutiny to disclosure requirements imposed on chatbots.¹⁵⁰

¹⁴⁶ S.B. 760, 103d Leg., Reg. Sess. (Mich. 2025).

¹⁴⁷ SB 482, 2026 Leg., Reg. Sess. (Fla. 2026). This one recently died with the end of the state’s legislative session, but it could pop up in a special session as soon as next month. See David McCabe & Patricia Mazzei, “How Trump Drove a Wedge Between Florida Republicans Over A.I.,” N.Y. Times (Mar. 16, 2026), [tinyurl.com/3tpktx36](https://www.nytimes.com/2026/03/16/us/politics/florida-republicans-ai.html).

¹⁴⁸ See *Nat’l Assoc. of Mfrs. v. SEC*, 800 F.3d 518, 522 (D.C. Cir. 2015) (holding that the loosened protection “is confined”—“emphatically”—“to advertising”).

¹⁴⁹ See *Am. Beverage Ass’n v. City & Cnty. of S.F.*, 916 F.3d 749, 763 (9th Cir. 2019) (Ikuta, J., concurring in the result). If you’re interested in this stuff, see me, “Social Media Transparency Rules, Zauderer Standard Head to Supreme Court,” note 119, above.

¹⁵⁰ Do not go saying something’s off here because AI firms are profit-seeking or because many people pay to use their services. You will look stupid. “The degree of First Amendment protection is not diminished merely

Lower First Amendment protection for disclosures is also supposed to be limited to the avoidance of deception.¹⁵¹ Again, this limit makes sense: it keeps the government from spraying disclosure requirements everywhere. Here again, some judges have drifted.¹⁵² A court faced with an “I’m not a human” disclosure requirement should wield some common sense. How stupid, it should ask, do you think kids are? Do they believe AI chatbots are human? Conceivably, a legislature could bar LLMs from actively deceiving people with false professions of corporeal status. But a disclosure requirement is infantilizing overkill—especially if required every hour.

Even if the lower First Amendment standard applies, Florida is not out of the woods. Fair enough, “I am not a human” is, in this context, a factual and uncontroversial statement. But isn’t requiring this disclosure unduly burdensome? Maybe one disclosure a week would be acceptable. But every chat? *Every hour*? Recall Europe’s effort to promote privacy by festooning the internet with irritating cookie popups. Florida is making a similar mistake.

SB 482—as well as California’s (enacted) SB 243 and the federal SAFE BOTs Act—require companion chatbots to periodically urge a minor to take a break.¹⁵³ To repeat: disclosure rules should trigger strict scrutiny outside the context of commercial transactions and the prevention of deception. There is also, here, the problem that *You should take a break* is not a cleanly factual statement. It is a viewpoint—indeed, *the government’s* viewpoint—being put into the mouth of the AI firm.¹⁵⁴

SB 243, the California law, requires an LLM to respond to any user who expresses thoughts of suicide or self-harm by referring her to a crisis hotline, and likely shutting down

because . . . speech is sold rather than given away.” *City of Lakewood v. Plain Dealer Publishing Co.*, 486 U.S. 750, 756 n.5 (1988); see also *303 Creative*, 603 U.S. at 594; *Sorrell*, 564 U.S. at 567. You pay to see a movie, after all, and “motion pictures are a significant medium for the communication of ideas.” *Burstyn*, 343 U.S. at 501. The mistake comes from thinking the AI output is *itself* the “commercial” speech. It is not—any more than a *New York Times* op-ed is “commercial” speech because you paid for the subscription. Commercial speech is speech *about* a commercial transaction. *Bd. of Trustees of State Univ. of N.Y. v. Fox*, 492 U.S. 469, 482 (1989).

¹⁵¹ One more time: details available at “Social Media Transparency Rules, Zauderer Standard Head to Supreme Court,” note 119, above.

¹⁵² See 916 F.3d at 767 (Nguyen, J., concurring in the judgment), as well as the dueling opinions in *Am. Meat Inst. v. U.S. Dep’t of Agric.*, 760 F.3d 18 (D.C. Cir. 2014) (en banc).

¹⁵³ S.B. 243, 2025–2026 Reg. Sess. (Cal. 2025); Safeguarding Adolescents From Exploitative BOTs Act, H.R. 6489, 119th Cong. (2025).

¹⁵⁴ Cf. *NetChoice v. Bonta*, 113 F.4th 1101, 1116-22 (9th Cir. 2024) (applying strict scrutiny to, and blocking, a requirement that websites compile and potentially disclose to the state reports on how they believe their content might harm minors).

the conversation.¹⁵⁵ The SAFE BOTs Act has a similar provision limited to minors. The referral mandate looks like another disclosure that's strongly laced with the government's opinion—*You are in danger of harming yourself and you should seek help*. That said, it's the kind of thing that could conceivably satisfy strict scrutiny, in that it serves a compelling interest (preventing suicide) through narrow means (meeting a person in crisis where she is). My problem is with the enforcement mechanism—a private right of action. This is an invitation to plaintiffs' lawyers to sue for every conceivable violation, sending AI firms into ass-covering mode. In theory, it's a nice law. In practice, all LLM users should expect, before long, to be getting crisis referrals at odd and annoying moments.

The refusal clause may be coming from a good place, but it's pernicious. A deeply troubled person turns to a chatbot for help or advice, and she's cut off mid-conversation and handed a phone number. This is naked censorship on the basis of content. It flies in the face of the right to information—the state is overruling the individual's choice about where to turn for support. It risks backfiring, severing discussion at a moment when doing so might be more dangerous than letting it continue. And the case for forced disengagement will only weaken with time, as AI gets ever better at handling such situations with care. The obvious, more narrowly tailored alternative is to require the crisis referral without silencing the LLM.

SB 482—back to Florida—bars minors from using a companion chatbot without a parent's consent. This returns us to Justice Scalia's observation that parental-consent mandates are really government mandates "subject only to parental veto."¹⁵⁶ As Scalia explained, such mandates make little sense. They subject third parties to heavy-handed state power on behalf of some parents, they work against the desires of many other parents (those fine with their kids seeing the pertinent speech), and they skip over the many ways parents can protect their children without state interference.¹⁵⁷ The less restrictive approach is for states to encourage the development and use of controls that let parents adjust settings, based on the age and maturity of their child, as they see fit.

Lurking behind all of these efforts is the problem of age verification. State-mandated requirements that apply to minors but not adults require, by implication, that a service figure

¹⁵⁵ SB 243, Sec. 22602(b)(1), requires a "protocol for preventing the production of suicidal ideation, suicide, or self-harm content to the user," as well as referrals "to crisis service providers" when the user "expresses suicidal ideation" or interest in self-harm. I suppose avoiding "suicidal ideation, suicide, or self-harm content" could be read to mean "The LLM can keep talking; it just can't make things worse." But that's meeting the legislative drafters more than halfway. They've made it so an AI firm's best move is to have its LLM clam up the moment there's a whiff of trouble.

¹⁵⁶ *Brown*, 564 U.S. at 795 n.3.

¹⁵⁷ *Id.* 801-04. And even if they did none of these things, such mandates would still ignore that minors have rights of their own. Recall once again Judge Posner's explanation as to why parents may not "enlist the aid of the state" in an effort to raise their kids in an "intellectual bubble." *Kendrick*, 244 F.3d at 576-77.

out who is a minor and who is an adult. Age verification thus places a “burden on the exercise” of First Amendment rights.¹⁵⁸ This is especially true online. Online age verification burdens users’ privacy, by forcing them to hand over sensitive information—IDs, biometric data, or both—to access lawful speech. It suppresses the speech of those who are unable, or understandably reluctant, to submit personal data before entering a website. It normalizes a “papers please” approach to society, habituating citizens to coughing up identification, and submitting to tracking and monitoring, wherever they go. Make no mistake, there is no clean technological fix here: every age-verification regime creates serious tradeoffs among privacy, security, accuracy, and access.¹⁵⁹

Any law that applies broadly, across all LLM interactions—laws that ban minors from using companion chatbots, or that require regular disclosures—carries with it an implicit age-verification requirement for all users who seek to use AI. This is an additional reason to conclude that such laws do not meet strict scrutiny (i.e., they are not the least restrictive means of tackling the problems they claim to address). Narrower laws—such as ones that shield minors from sexually explicit chats or suggestions of suicide—should ensure that any age-verification burden sits at the source of the potential harm. The LLM shouldn’t ask you to provide proof of your age unless you enter a prompt that will call forth material minors have no constitutional right to see.

Right as I was ready to shove this thing out the door, the White House released a national AI policy framework.¹⁶⁰ It endorses, among other things, “robust tools” for parents and guardians and “commercially reasonable” age-assurance requirements. Parental controls are great, and age assurance suggests something better than age verification.¹⁶¹ The devil, though, as usual, is in the details. It is constitutionally safer for Congress to promote parental controls than to mandate them. Letting parents set even strict limits on minors’ accounts is likely legal—and, anyway, something AI firms can live with—but there may be some outer

¹⁵⁸ *Free Speech Coalition*, 145 S. Ct. at 2309.

¹⁵⁹ This paragraph is a riff on Eric Goldman, *The “Segregate-and-Suppress” Approach to Regulating Child Safety Online*, 28 Stan. Tech. L. Rev. 173 (2025). Which is not to say I lack thoughts of my own on the matter. See Corbin K. Barthold, “Age-Verification Laws Are a Verified Mistake,” *Law & Liberty* (Jan. 9, 2025), [tinyurl.com/mr26r2d3](https://www.tinyurl.com/mr26r2d3).

¹⁶⁰ The White House, *National Policy Framework for Artificial Intelligence: Legislative Recommendations* (Mar. 20, 2026), [tinyurl.com/3k3vp9ps](https://www.tinyurl.com/3k3vp9ps).

¹⁶¹ The difference between “age verification” and “age assurance” is far from settled. But I take the White House’s use of “assurance” as a gesture at options less onerous than “verification.” Assurance is widely understood as an umbrella term that includes the possibility of estimation. Verification ≈ *We think we know your age: you showed your ID*. Estimation ≈ *We have a sense of your age: your behavior on our platform / signals related to your email address / your use of a credit card suggests you’re an adult*. (But notice how, even for verification, I use the word “think.” All methods can be gamed or evaded by a minor who is clever and motivated—as most of them are. Cf. Byron Kaye, “One-Fifth of Australian Teens Still Use TikTok, Snapchat After Social Media Ban,” *Reuters* (Mar. 12, 2026), [tinyurl.com/4styp5mr](https://www.tinyurl.com/4styp5mr).)

edge where such controls rub up against teens' independent First Amendment rights.¹⁶² Age assurance can be done in a number of ways, some less intrusive and problematic than others. And there is again my point about placing barriers in precise spots, rather than at the entrance to the entire service. Always, a court should ask whether the measure in question functions as a burden likely to chill expression or block lawful access to information. If it does, strict scrutiny probably applies, and it is up to the government to show that it has not overshot the mark. The First Amendment demands no less—and if I've done my job, you see what I mean.

¹⁶² See Berin Szóka, "Congressional Republicans Push Bills That Would Block Kids' Access to Content for Ideological Reasons," Techdirt (Mar. 9, 2026), tinyurl.com/2wzh63k8.