NONHUMAN SPEAKERS AND THE CONSTITUTION

*Helen Norton**

Long-standing free speech law sometimes does—and sometimes does not—turn on speakers' human character. This Essay explores how these doctrinal threads might shape the constitutional doctrine applied to generative artificial intelligence's (AI) speech outputs if we assume that a human speaker cannot be identified as the creator of those words and images—and that those words and images are nevertheless considered "speech" for First Amendment purposes.

For example, First Amendment doctrine sometimes requires proof of a speaker's culpable mental state before that speaker can be held liable for what would otherwise be unprotected speech. These requirements seek to prevent the chilling of human speakers' valuable expression and to ensure the punishment only of morally blameworthy human speakers. But because generative AI's speech outputs do not pose these concerns, these doctrinal features requiring a human speaker should not apply to the government's regulation of generative AI's outputs—lest those outputs be insulated from liability for endangering human users.

Other long-standing First Amendment doctrines apply to the government's regulation of speech regardless of the speaker's human or nonhuman character. These doctrinal threads show that extending First Amendment coverage to nonhuman speakers' expression need not foreclose the government's regulatory efforts to protect human users from grave harm. At the same time, these doctrinal threads also show that First Amendment coverage of the words and images produced by generative AI can provide important guardrails against the government's self-interested or clumsy regulation of those outputs despite their value to human users.

---

* University Distinguished Professor and Rothgerber Chair in Constitutional Law, University of Colorado School of Law. Thanks to the *Fordham Law Review*, Alexander Tsesis, and Caroline Mala Corbin for their work in convening this symposium. Thanks also for thoughtful questions and comments from the symposium's participants, the Rocky Mountain First Amendment All Stars, participants at Yale Law School's Free Expression Scholars Conference, and from RonNell Andersen Jones, Danielle Keats Citron, and Toni M. Massaro. Finally, thanks to Jonah Dixon, Armania Heckenmueller, and Sarah Sorenson for excellent research assistance.

This Essay was prepared for the Symposium titled *Constitutional Law in the Age of Generative AI* hosted by the *Fordham Law Review* on February 27, 2026, at Fordham University School of Law.

Finally, like its free speech doctrine more generally, the U.S. Supreme Court's government speech doctrine generally remains indifferent to governmental speakers' human or nonhuman character. Here, too, governments' increasing use of generative AI for expressive purposes calls for an appropriately listener-centered government speech doctrine: one that recognizes the inevitability (and often the value) of government speech while constraining a limited but important universe of government speech that interferes with its targets' exercise of their constitutional rights.

INTRODUCTION 526

I. FIRST AMENDMENT DOCTRINAL FEATURES THAT REQUIRE A HUMAN SPEAKER 529

II. FIRST AMENDMENT DOCTRINAL FEATURES THAT DO NOT REQUIRE A HUMAN SPEAKER 533

III. USING INTERMEDIATE SCRUTINY TO OPERATIONALIZE AN APPROPRIATELY LISTENER-CENTERED DOCTRINE 539

IV. LISTENER-CENTERED GOVERNMENT SPEECH DOCTRINE 544

CONCLUSION 550

INTRODUCTION

"Generative AI" refers to artificial intelligence (AI) systems that learn from statistical patterns revealed by large datasets (i.e., large language models (LLMs)) to generate new content in the form of words and images that resemble human-created material.[1] This Essay considers questions regarding the constitutional doctrine to be applied to the government's regulation of the words and images produced by generative AI *if* we assume that (1) a human speaker cannot be identified as the creator of those words and images[2] and (2) those words and images are nevertheless considered "speech" for First Amendment purposes.[3]

Of course, those assumptions are contestable and contested. Many thoughtful scholars, lawyers, and judges have wrestled, and continue to wrestle, with important threshold questions of whether these sorts of outputs should be understood as speech covered by the First Amendment. Some conclude that these outputs cannot be understood as the expression of a human speaker, and thus cannot be treated as speech for First Amendment

1. *See* Harry Surden, *ChatGPT, AI Large Language Models, and Law*, 92 FORDHAM L. REV. 1941, 1957 (2024).

2. *See* Moody v. NetChoice, LLC, 144 S. Ct. 2383, 2410 (2024) (Barrett, J., concurring) (anticipating circumstances where humans "hand the reins to an AI tool").

3. *See* 303 Creative LLC v. Elenis, 143 S. Ct. 2298, 2312 (2023) (explaining that "[a]ll manner of speech—from 'pictures, films, paintings, drawings, and engravings,' to 'oral utterance and the printed word'—qualify for the First Amendment's protections" (quoting Kaplan v. California, 413 U.S. 115, 119–20 (1973))).

purposes.[4] Others propose that these outputs *can* be traced to human creators, and thus should be treated as speech covered by the First Amendment.[5] Still other possibilities exist, as Professors Margot Kaminski and Meg Leta Jones remind us, because a First Amendment lens is by no means the only (and perhaps not the best) perspective for thinking about the potential regulation of generative AI speech outputs.[6]

Engaging the question of whether and when the First Amendment covers AI outputs is hard and important work, and I applaud those engaged in these various inquiries.[7] But *if* we choose to consider these questions through a First Amendment lens, I suggest that long-standing Free Speech Clause doctrine offers guidance for assessing the government's regulation of the words and images produced by generative AI *even if* a human speaker cannot be identified as the creator of those words and images *and* those outputs are nevertheless treated as "speech" for First Amendment purposes.

Although that combination of assumptions is neither inevitable nor necessarily desirable, neither is it implausible. In earlier work, Professors Toni M. Massaro, Margot Kaminski, and I concluded that the Free Speech Clause can reasonably be understood to apply to AI outputs even if those outputs cannot be ascribed to a human speaker: This is because First Amendment law so often seeks to safeguard listeners' interests and constrain the government's potential for overreach rather than protect speakers as speakers.[8] Descriptive and conceptual, that work considered what could be,

4. *See, e.g.*, Mackenzie Austin & Max Levy, *Speech Certainty: Algorithmic Speech and the Limits of the First Amendment*, 77 STAN. L. REV. 1, 11, 67–83 (2025) (asserting that the First Amendment does not protect the outputs of machine learning algorithms because "[s]peech is protected only when the speaker knows with certainty what he says at the moment he says it"—knowledge unavailable to AI outputs); Peter N. Salib, *AI Outputs Are Not Protected Speech*, 102 WASH. U. L. REV. 83, 112–34 (2024) (asserting that generative AI outputs do not reflect communication traceable to a human creator and thus are not speech protected by the First Amendment). *See generally* Alexander Tsesis, *Algorithmic Models and Human Free Speech*, 95 FORDHAM L. REV. 599 (2026) (arguing that machine-generated data should not be considered speech under the First Amendment but rather artificial language tokens).

5. *See, e.g.*, Francesca Procaccini & Wendy K. Tam, *Algorithmic Speech*, 115 CALIF. L. REV. (forthcoming 2027) (manuscript at 1) (on file with the *Fordham Law Review*) (asserting that AI outputs are covered by the First Amendment because they "embed human judgment, values, and expressive intent at every stage of their design and deployment").

6. *See* Margot E. Kaminski & Meg Leta Jones, *Constructing AI Speech*, 133 YALE L.J.F. 1212, 1221–22 (2024) (explaining that generative AI outputs could instead be understood, among other possibilities, as tools, data processors, or complex risky systems governed by tort or consumer protection law).

7. *See id.* at 1265 ("A stable legal construction of AI has not yet set in, and negotiations are occurring with all the histories of global and national technology governance as the setting.").

8. *See* Toni M. Massaro & Helen Norton, *Siri-ously?: Free Speech Rights and Artificial Intelligence*, 110 NW. U. L. REV. 1169, 1172–74 (2016); Toni M. Massaro, Helen Norton & Margot E. Kaminski, *SIRI-OUSLY 2.0: What Artificial Intelligence Reveals About the First Amendment*, 101 MINN. L. REV. 2481, 2486–96 (2017).

what might be, what perhaps will be—but did not insist on what should be or what must be.[9]

Here, I examine threads of traditional First Amendment doctrine that apply to speech valuable to human listeners regardless of the speaker's human or nonhuman character. And these doctrinal threads, I posit, show that determining nonhuman speakers' expression to be covered by the First Amendment need not insulate that expression from the government's regulation to protect human users from grave harm, notwithstanding some observers' worries.[10] At the same time, First Amendment coverage of the words and images produced by generative AI—even if those outputs are untraceable to a human speaker—provides important guardrails against the government's self-interested or clumsy regulation of those outputs to the detriment of human users.

That nonhuman speakers might produce words and images of value (and sometimes of danger) to human listeners has implications not only for the doctrine that applies to government's regulation of generative AI, but also for the doctrine that applies to the government's own use of generative AI to develop and deliver expression. The U.S. Supreme Court's government speech doctrine is similarly inattentive to government speakers' human character.[11] Governments' increasing use of generative AI for expressive purposes calls for an appropriately listener-centered government speech doctrine: one that recognizes the inevitability (and often the value) of government speech while constraining a limited but important universe of government speech that interferes with its targets' exercise of their constitutional rights.

A few caveats before we continue. First, this Essay makes no claims about how generative AI works now or how it might work at some point in the future.[12] Nor do I propose any specific regulatory intervention to protect the

9. *See* Massaro & Norton, *supra* note 8, at 1172–73; Massaro, Norton & Kaminski, *supra* note 8, at 2524; *see also id.* at 2485 (acknowledging that "the logic of free speech law and doctrine can carry us to unfamiliar and uncomfortable, or even dangerous, places").

10. *See* Salib, *supra* note 4, at 86, 88 (predicting that "[i]f AI outputs are best understood as protected speech, then laws regulating them directly, even to promote safety, will have to satisfy the strictest constitutional tests" and that the resulting "consequences for lawmaking—and indeed, for humanity—could be dire"). These concerns that the First Amendment threatens to strip the government of its power to address serious harm posed by generative AI outputs are exacerbated by the contemporary U.S. Supreme Court's methodology for identifying categories of unprotected speech, which now requires the government to show that the regulation of speech within a contested category has been historically treated as exempt from First Amendment review—a showing nearly impossible to make with respect to many new expressive technologies. *See* United States v. Stevens, 559 U.S. 460, 471–72 (2010); *see also* Steven H. Shiffrin, *The Dark Side of the First Amendment*, 61 UCLA L. REV. 1480, 1489, 1493 (2014) (describing this as the "frozen categories" approach).

11. *See infra* notes 105–13 and accompanying text.

12. *See generally* Harry Surden, *Artificial Intelligence and Law—An Overview of Recent Technological Changes: Keynote Address at the 2024 IRA C. Rothgerber Jr. & Silicon Flatirons Conference on Artificial Intelligence and Constitutional Law*, 96 U. COLO. L. REV. 375 (2025) (explaining the technology at work in generative AI); Lyrissa Lidsky & Andrew Daves, *Inevitable Errors: Defamation by Hallucination in AI Reasoning Models*, 6 J. FREE SPEECH L. 477, 517–23 (2025) (similar).

human users of generative AI from its harms. Instead, I muse about the free speech doctrine to be applied to those regulatory interventions, whatever they may be. Finally, this Essay doesn't address the First Amendment analysis applicable to government's regulation of humans' use of generative AI outputs to inform, produce, or distribute their own speech.[13] Instead, I consider whether and when the First Amendment allows the government's regulation of generative AI *as a producer of words and images itself*—as is the case, for example, of a chatbot's response to a human's prompt.[14]

This Essay unfolds as follows: Part I sketches First Amendment doctrinal features that *require* a human speaker and then asserts that those features should not apply when assessing the constitutionality of the government's regulation of generative AI's speech outputs. Part II then discusses Free Speech Clause doctrine that *doesn't* turn on speakers' human character, and how those doctrinal threads might inform the doctrine applied to the regulation of generative AI speech outputs. Part III next explores, in a bit more detail, how intermediate scrutiny can operationalize an appropriately listener-centered doctrine for evaluating the government's regulation of generative AI speech outputs. Finally, Part IV pivots to consider how a listener-centered government speech doctrine might address the constitutional questions raised by the government's deployment of generative AI for its own expressive purposes.

I. FIRST AMENDMENT DOCTRINAL FEATURES THAT REQUIRE A HUMAN SPEAKER

Some features of traditional First Amendment doctrine require a human speaker, and some don't.[15] For example, First Amendment doctrine sometimes requires proof of the speaker's culpable mental state before that speaker can be held liable for what would otherwise be unprotected speech. In many settings, defamatory falsehoods don't trigger liability absent proof of the speaker's actual malice.[16] Nor can a speaker be held liable for threatening statements that cause a reasonable person to fear for her physical safety absent proof of the speaker's conscious disregard of the risk that their

13. To illustrate, consider a human user who uses generative AI to convert a photograph of a dressed person into a nude portrayal of that person and then publishes that depiction without the subject's consent. There, the expression can be directly traced to a human, and that human's mental state may be relevant to the First Amendment analysis applicable to the government's regulation of that human's expression. This Essay does not address those very different First Amendment questions.

14. *See, e.g.*, *infra* notes 116–18 and accompanying text.

15. Conversely, the First Amendment can protect humans' speech even absent any human listener. Consider a human singing in the shower or keeping a private journal with the intention that no other human will ever hear or see that expression. *See* C. EDWIN BAKER, HUMAN LIBERTY AND FREEDOM OF SPEECH 51 (1989) (conceptualizing free speech doctrine to protect "solitary" speech like "keeping a diary" or "singing"). Think, too, of a human's speech to a chatbot companion. *See* Marc Jonathan Blitz, XR, AI, and the Doctrinal Distinctiveness of Solitary Speech 24–34 (2026) (unpublished manuscript) (on file with the *Fordham Law Review*) (describing humans' solitary interactions with AI as part of their freedom of thought potentially protected by the First Amendment).

16. *See* N.Y. Times Co. v. Sullivan, 376 U.S. 254, 279–80 (1964).

statement would cause such fear.[17] And the First Amendment doesn't allow the government to punish speech that incites violence or illegal activity without evidence that the speaker intended that result.[18]

Requiring proof of the speaker's culpable mental state in these settings seeks to prevent the chilling of speakers' valuable expression and to protect those speakers' exercise of their individual autonomy. For instance, the Supreme Court explained its decision to require a showing of the defendant's reckless mental state as a condition for punishing true threats by detailing how the prospect of governmental punishment can chill a speaker's expression: "The speaker's fear of mistaking whether a statement is a threat; his fear of the legal system getting that judgment wrong; his fear, in any event, of incurring legal costs—all those may lead him to swallow words that are in fact not true threats."[19] Requiring this proof also addresses deontic concerns by imposing punishment only on speakers with a morally blameworthy mental state.[20]

But these concerns about chilling, autonomy, and moral blameworthiness assume a human speaker who may be chilled into silence by the prospect of the government's punishment, who makes her own autonomous choices, and who exercises moral agency. Because generative AI doesn't have a mental state, culpable or otherwise, generative AI speech outputs don't pose those concerns.[21] And because generative AI has no mental state, of course its mental state cannot be proven.

Insisting on (unobtainable) evidence of generative AI's culpable mental state as a prerequisite for establishing liability would thus provide those outputs with more First Amendment protection than that provided to human speech—even though those outputs are no more valuable (and perhaps are less valuable) and no less harmful (and perhaps are more harmful) to their human listeners than those produced by human speakers.[22] Applying such doctrinal features to the government's regulation of generative AI speech outputs that cause harm to human users would deny relief from very real

17. *See* Counterman v. Colorado, 143 S. Ct. 2106, 2111–12 (2023).

18. *See* Brandenburg v. Ohio, 395 U.S. 444, 447 (1969) (per curiam); *see also Counterman*, 143 S. Ct. at 2115–16 (describing the *Brandenburg* rule).

19. *Counterman*, 143 S. Ct. at 2116; *see also N.Y. Times*, 376 U.S. at 271–72 (explaining that the First Amendment protects defamatory speech critical of public officials absent evidence of the speaker's actual malice to ensure breathing room for valuable expression).

20. *See Counterman*, 143 S. Ct. at 2118 (describing defendants with a "reckless" mental state as "morally culpable" because they "have done more than make a bad mistake. They have consciously accepted a substantial risk of inflicting serious harm.").

21. *See* Kaminski & Jones, *supra* note 6, at 1240 ("AI systems do not themselves experience a chilling effect. AI-generated speech, like commercial speech, could be characterized as inherently hardier than ordinary speech.").

22. *See* Nina Brown, *Bots Behaving Badly: A Products Liability Approach to Chatbot-Generated Defamation*, 3 J. FREE SPEECH L. 389, 401 (2023) ("Such an outcome would give generative AI systems a pass for causing these harms, and any other harms . . . where there is a mens rea requirement. This would be unsound because it would remove incentives for developers to invest in products that minimize speech harms. (It also does not seem a just result when the harm a plaintiff suffers may be every bit as real when the speaker is a chatbot as when it is a human.)").

injuries for no reason other than an unconsidered commitment to existing doctrine. Defamatory falsehoods, for instance, inflict reputational damage upon their subjects regardless of the speaker's mental state.[23] So, too, do "true threats" inflict life-disrupting (including speech-silencing) harm, regardless of the speaker's mental state, by causing their target reasonably to fear for her physical safety.[24]

For these reasons, doctrinal features that require a human speaker should not apply when assessing First Amendment challenges to the government's regulation of the speech outputs produced by nonhuman speakers like generative AI. This may seem obvious, but recent experience indicates otherwise, as at least some litigants and judges try to force the square-peg features of traditional free speech doctrine that require a human speaker into the round hole of generative AI speech outputs. For instance, the first court to consider a defamation claim involving generative AI speech outputs insisted on finding a place for the "actual malice" requirement even absent a human speaker capable of a malicious mental state.[25] And, facing accusations that its chatbot characters' outputs caused young users to harm themselves, a tech company argued that those outputs were not regulable as unprotected incitement because generative AI could not "intend" to incite unlawful conduct.[26]

To be fair, courts' and litigants' reflexive deployment of long-standing doctrine developed to protect human speakers even when those human speakers are absent may reflect their uncertainty about what doctrinal features *should* apply under those circumstances. One possibility would be to apply the traditional doctrine without its human-centered mental state elements, thus establishing liability for harm caused by AI outputs that otherwise satisfy the elements of unprotected threats, defamation, or

---

23. *See* Gertz v. Robert Welch, Inc., 418 U.S. 323, 347–50 (1974) (finding that the First Amendment poses no bar to a private figure's recovery of actual damages for the reputational harm inflicted by defamatory falsehoods on matters of public concern, even absent a showing of the speaker's actual malice); Dun & Bradstreet, Inc. v. Greenmoss Builders, Inc., 472 U.S. 749, 761 (1985) (plurality opinion) (holding that the First Amendment poses no bar to a private figure's recovery of presumed and punitive damages for defamatory falsehoods on matters of private concern even absent a showing of actual malice because of the reputational harm inflicted by those falsehoods).

24. *See Counterman*, 143 S. Ct. at 2113 (explaining that a statement can count as a true threat "based solely on its objective content"); *see also* Helen Norton, *Fear and Free Speech*, 98 S. CAL. L. REV. 1351, 1358–62 (2025) (detailing the ways in which true threats silence their targets' protected expression).

25. *See* Walters v. OpenAI, LLC, No. 23-A-04860-2, slip op. at 16–18 (Ga. Super. Ct. May 19, 2025) (requiring the plaintiff to demonstrate that a specific individual at OpenAI knew the challenged output was false for purpose of establishing actual malice). As Professor Nina Iacono Brown observes, "[r]equiring a plaintiff to prove the developer had a particular mens rea with respect to the subject text of the defamation claim makes no more sense than asking it to prove that the chatbot itself had a particular mens rea." Brown, *supra* note 22, at 401.

26. *See* Character Techs., Inc.'s Motion to Dismiss Plaintiff's First Amended Complaint (Doc. 11) at 12–13, Garcia v. Character Techs., Inc., 785 F. Supp. 3d 1157 (M.D. Fla. 2025) (No. 24-cv-01903) [hereinafter Character Techs., Inc.'s Motion to Dismiss] ("There are no allegations that [Character AI] intended any lawless or violent action.").

incitement. In other words, when there's no human speaker with First Amendment interests to protect, there's no reason to protect these statements once the claim's other elements are proven. For defamation purposes, this would mean that the plaintiff must prove that a chatbot published a reputation-damaging falsehood about her—but not that the nonhuman speaker had any particular mental state.[27] For purposes of identifying true threats, this would require the prosecutor or plaintiff to prove that a reasonable person in the target's position would experience the defendant's statements to threaten her physical safety—but not that the generative AI that produced those statements had any particular mental state.[28] And for incitement, this would require the prosecutor to show that the statement was substantially likely to incite imminent violence or other illegal action, period.[29]

Even so, application of traditional doctrinal requirements that don't require a human speaker might, in some circumstances, support the conclusion that AI's nonhuman character defeats an entirely different element necessary to establish liability. Professor Lyrissa Barnett Lidsky and Andrew Daves, for instance, recognize the inappropriateness of applying the "actual malice" standard to generative AI's defamatory falsehoods given its nonhuman character.[30] They nevertheless propose that generative AI is better understood as a "distributor" (rather than a "publisher") of such falsehoods for purposes of defamation law and thus insulated from liability.[31]

Still other possibilities call for applying altogether different doctrinal standards to generative AI outputs. As Kaminski and Jones observe, areas of the law other than the First Amendment—like product liability and consumer protection law—"construct AI systems very differently," and without attention to AI's lack of human intent.[32] Along these lines, Professors Ian Ayres and Jack M. Balkin describe the law of AI as "the law of risky agents without intentions," in which such risky agents' outputs are assessed by objective standards rather than by standards that require a showing of an actor's culpable mental state: "Holding AI agents to objective standards of behavior, in turn, means holding the people and organizations that implement these technologies to standards of reasonable care and requirements of reasonable reduction of risk."[33] These various possibilities share common ground in their (appropriate, in my view) resistance to applying First

27. *See* RESTATEMENT (SECOND) OF TORTS § 558 (A.L.I. 1977) (outlining the general elements of defamation).

28. *See Counterman*, 143 S. Ct. at 2113.

29. *See* Brandenburg v. Ohio, 395 U.S. 444, 447 (1969) (per curiam).

30. *See* Lidsky & Daves, *supra* note 12, at 497–506.

31. *See id.* (suggesting that generative AI be treated as a distributor and not a publisher for purposes of defamation law); *see also id.* at 542 (noting that over time, the common law of defamation has evolved "to ensure that those who provide public access to the newest communications tools" do not face devastating liability).

32. Kaminski & Jones, *supra* note 6, at 1212.

33. Ian Ayres & Jack M. Balkin, *The Law of AI is the Law of Risky Agents Without Intentions*, U. CHI. L. REV. ONLINE (Nov. 27, 2024), https://lawreview.uchicago.edu/online-archive/law-ai-law-risky-agents-without-intentions [https://perma.cc/4VMT-U9TD].

Amendment doctrinal features requiring a human speaker to the government's regulation of generative AI outputs.

II. FIRST AMENDMENT DOCTRINAL FEATURES THAT DO NOT REQUIRE A HUMAN SPEAKER

Next, we turn to long-standing doctrines that assess First Amendment challenges to the government's regulation of what courts and scholars agree is "speech" *without* imposing any features that assume, much less require, a human speaker. This is most notably the case of the commercial speech doctrine, which focuses almost entirely on commercial expression's value or danger to its human listeners rather than on the speaker's human or nonhuman character.[34] And those doctrinal traditions, in turn, illuminate doctrinal possibilities for the regulation of generative AI's speech outputs.

The Supreme Court's commercial speech doctrine, at least as initially articulated, emphasized listeners' (i.e., consumers') reliance on accurate commercial speech to inform their First Amendment interests in autonomous decision-making. Recall *Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council, Inc.*,[35] where the Supreme Court first made clear that commercial speech is worthy of some level of First Amendment protection.[36] There, the Court struck down Virginia's law barring pharmacists from advertising the price of any prescription drug, thus effectively denying this information's availability to the public.[37] The Court's opinion focused on the forbidden expression's informational value to the consumers of prescription drugs.[38] As the majority explained:

> As to the particular consumer's interest in the free flow of commercial information, that interest may be as keen, if not keener by far, than his interest in the day's most urgent political debate. Appellees' case in this respect is a convincing one. Those whom the suppression of prescription drug price information hits the hardest are the poor, the sick, and particularly the aged. A disproportionate amount of their income tends to be spent on prescription drugs; yet they are the least able to learn, by shopping from pharmacist to pharmacist, where their scarce dollars are best spent. When drug prices vary as strikingly as they do, information as to who is charging what becomes more than a convenience. It could mean the alleviation of physical pain or the enjoyment of basic necessities.[39]

34. *See infra* notes 57–60 and accompanying text.

35. 425 U.S. 748 (1976).

36. *See id.* at 761–70 (finding that commercial speech receives some level of First Amendment protections).

37. *See id.* at 751–54.

38. *See id.* at 753–54.

39. *Id.* at 763–64; *see also id.* at 754 (noting that the price of prescription drugs varied by as much as 650 percent even in the same locality—information made effectively unavailable to consumers by the state's ban); Alan B. Morrison, *How We Got the Commercial Speech Doctrine: An Originalist's Recollections*, 54 CASE W. RSRV. L. REV. 1189, 1192 (2004) ("[T]hese laws all suppressed truthful and useful information, and at their core they seemed to run directly counter to the First Amendment.").

The Court then turned to commercial expression's value to the public more generally, observing that commercial speech can advance listeners' democratic self-governance and enlightenment interests in addition to their autonomy interests: "Generalizing, society also may have a strong interest in the free flow of commercial information. Even an individual advertisement, though entirely 'commercial,' may be of general public interest."[40]

The Court did not, however, engage with the regulated commercial speakers' (in that case, human pharmacists') expressive interests. Instead, it expressly assumed that those speakers' economic self-interest leaves their speech less susceptible to potential chilling by the government's regulation.[41] The Court also indicated that a lower level of First Amendment protection is appropriate in this setting because commercial expression's accuracy "may be more easily verifiable by its disseminator" due to the commercial speaker's greater knowledge about the products or services "that he himself provides and presumably knows more about than anyone else."[42]

A few years later, the Court operationalized these intuitions through a doctrinal framework that made clear that commercial expression's First Amendment protection turns on its value—or instead its harm—to listeners.[43] In *Central Hudson Gas & Electric Corp. v. Public Service Commission of New York*,[44] the Court considered an electric utility company's First Amendment challenge to the government's ban on advertising that promoted consumers' energy use.[45] The Court announced a form of intermediate scrutiny, requiring the government to show that its (even content-based) regulation of commercial speech "directly advances" its substantial interest in a way "not more extensive than is necessary to serve that interest."[46] In invalidating the ban, the Court emphasized that the contested expression served human listeners' interests in "the fullest possible dissemination of information."[47]

At the same time, the Court's newly articulated doctrine emphasized that commercial speech that undermines listeners' First Amendment interests receives no First Amendment protection, thus enabling the government's content-based prohibitions on such speech:

---

40. *Virginia Pharmacy*, 425 U.S. at 764.

41. *See id.* at 771 n.24 (observing that commercial speakers' profit motivations may mean that commercial speech is "more durable" than other speech); *see also* Dun & Bradstreet, Inc. v. Greenmoss Builders, Inc., 472 U.S. 749, 762 (1985) (plurality opinion) (describing commercial speech that is "solely motivated by the desire for profit" as "hardy" and "less likely to be deterred than" other speech).

42. *Virginia Pharmacy*, 425 U.S. at 771 n.24.

43. *See* Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n, 447 U.S. 557, 563 (1980).

44. 447 U.S. 557 (1980).

45. *See id.* at 559–60.

46. *Id.* at 566.

47. *Id.* at 561–62; *see also id.* at 563 ("The protection available for particular commercial expression turns on the nature both of the expression and of the governmental interests served by its regulation.").

> The First Amendment's concern for commercial speech is based on the informational function of advertising. Consequently, there can be no constitutional objection to the suppression of commercial messages that do not accurately inform the public about lawful activity. The government may ban forms of communication more likely to deceive the public than to inform it, or commercial speech related to illegal activity.[48]

This doctrinal framework does not include features that turn on the speaker's human character, as would have been the case had the doctrine required a showing of the speaker's subjective mental state. Instead, false or misleading commercial speech receives no First Amendment protection—regardless of the commercial speaker's mental state—because the harm to listeners' informed decision-making does not turn on the speaker's mental state.[49] In other words, consumers experience harm from false or misleading commercial speech regardless of the speaker's human or nonhuman character.

At the same time, the Court also concluded that commercial speech that is not false, misleading, or related to illegal activity is generally valuable to listeners, and for that reason its regulation triggers intermediate scrutiny.[50] But why not strict scrutiny? In other words, why did the Court provide "less protection to commercial speech than to other constitutionally safeguarded forms of expression"?[51] As Professor Ashutosh Bhagwat explains, courts often choose a form of intermediate scrutiny in settings where they have concluded that strict scrutiny should *not* apply because of the contested expression's lower First Amendment value, its potential threat to weighty public interests, and the absence of reasons to distrust the government's regulatory capacity in that setting.[52] All of these factors are in play in the

48. *Id.* at 563–64 (internal citations omitted); *see also* Zauderer v. Off. of Disciplinary Couns., 471 U.S. 626, 638 (1985) ("The States and the Federal Government are free to prevent the dissemination of commercial speech that is false, deceptive, or misleading."). For related reasons, the commercial speech doctrine provides the government with substantial leeway to require commercial speakers to make truthful disclosures about their goods and services because such disclosures so often serve listeners' informational interests. *See id.* at 651–52.

49. *See, e.g.*, Omnicare, Inc. v. Laborers Dist. Council Constr. Indus. Pension Fund, 575 U.S. 175, 179 (2015) (explaining that a buyer "need not prove . . . that the defendant acted with any intent to deceive" in establishing a violation of Section 11 of the Securities Act of 1933); Aaron v. SEC, 446 U.S. 680, 696–702 (1980) (explaining that Section 17(a)(3) of the Securities Act of 1933 "quite plainly focuses upon the *effect* of particular conduct on members of the investing public, rather than upon the culpability of the person responsible"); FTC v. Algoma Lumber Co., 291 U.S. 67, 81 (1934) (explaining that advertisers' "innocence of motive" does not relieve them of liability for violating the Federal Trade Commission Act's bar on deceptive trade practices).

50. *See Cent. Hudson*, 447 U.S. at 564.

51. Bolger v. Youngs Drug Prods. Corp., 463 U.S. 60, 64–65 (1983).

52. *See* Ashutosh Bhagwat, *The Test That Ate Everything: Intermediate Scrutiny in First Amendment Jurisprudence*, 2007 U. ILL. L. REV. 783, 824; *see also Bolger*, 463 U.S. at 65 (explaining that content-based restrictions on commercial speech are permissible because "of the greater potential for deception or confusion" in that context).

commercial speech context.[53] The Court also explained its concern that requiring strict scrutiny of the government's regulation of commercial speech might lead courts to dilute that review in order to uphold sensible regulatory frameworks—a choice that would endanger strict scrutiny's vigorous protection of speech in other settings.[54]

To be sure, the Court's later commercial speech decisions reflected a different tone, with less focus on commercial expression's hardiness (and thus its resistance to chilling) and more on the dangers of government paternalism.[55] But concerns about governmental paternalism are themselves listener centered, steeped as they are in worries that the government's interference with listeners' interests denies those listeners access to valuable information.[56]

I don't claim that the Court that created the commercial speech doctrine understood itself to be engaging a First Amendment question involving—much less developing a doctrine for—nonhuman speakers.[57] Instead, my point is that the *Central Hudson* Court's doctrinal intuitions centered solely on listeners' interests such that the speaker's human or nonhuman character was immaterial. And these listener-centered intuitions provide a template for

53. *See* Ohralik v. Ohio State Bar Ass'n, 436 U.S. 447, 455–56 (1978) (describing the government's "traditional[]" regulation of commercial expression and such expression's "subordinate" First Amendment value).

54. *See id.* ("We have not discarded the 'common-sense' distinction between speech proposing a commercial transaction, which occurs in an area traditionally subject to government regulation, and other varieties of speech. To require a parity of constitutional protection for commercial and noncommercial speech alike could invite dilution, simply by a leveling process, of the force of the Amendment's guarantee with respect to the latter kind of speech. Rather than subject the First Amendment to such a devitalization, we instead have afforded commercial speech a limited measure of protection, commensurate with its subordinate position in the scale of First Amendment values, while allowing modes of regulation that might be impermissible in the realm of noncommercial expression." (internal citations omitted)).

55. *See, e.g.*, 44 Liquormart, Inc. v. Rhode Island, 517 U.S. 484, 503 (1996) (plurality opinion) ("The First Amendment directs us to be especially skeptical of regulations that seek to keep people in the dark for what the government perceives to be their own good."); *see also* Bhagwat, *supra* note 52, at 794–95, 829 (describing this antipaternalistic, and less government-friendly, turn in commercial speech doctrine).

56. *See* Helen Norton, *Manipulation and the First Amendment*, 30 WM. & MARY BILL RTS. J. 221, 242 (2021) ("[P]aternalism describes interference with autonomous choices that others think unwise, rather than interference with practices that themselves frustrate autonomous choice.").

57. Nor do I claim that the contemporary Court remains committed to its earlier approach to commercial speech. *See* Sorrell v. IMS Health, Inc., 564 U.S. 552, 564–66, 571–72 (2011) (noting the possibility of applying strict scrutiny to the government's regulation of commercial speech while holding that the contested regulation fails even *Central Hudson* intermediate scrutiny). But lower courts have so far largely remained reluctant to retreat from the Court's long-standing doctrine. *See* William D. Araiza, *Invasion of the Content-Neutrality Rule*, 2019 BYU L. REV. 875, 912 ("[E]ven relatively recent lower court opinions have continued to resist imposing strict scrutiny on commercial speech regulations, despite *Sorrell*'s implication that they should."); *see also* Alan B. Morrison, *No Regrets (Almost): After* Virginia Board of Pharmacy, 25 WM. & MARY BILL RTS. J. 949, 958–60 (2017) (criticizing some of the Court's more recent commercial speech decisions as insufficiently listener-centered when those decisions did not rest on any "concern that the recipients of the information at issue are kept in the dark against their will").

considering the First Amendment questions raised by the government's regulation of the speech outputs of nonhuman speakers like generative AI.

Commercial speech doctrine is among the most prominent illustrations of a listener-centered doctrine that attends little, if at all, to the speaker's First Amendment interests—and thus attends little, if at all, to the speaker's human or nonhuman character. But it's not the only illustration. In *Citizens United v. Federal Election Commission*,[58] a very different context involving speech in public discourse, the Court emphasized the First Amendment value of corporate speech not to corporate speakers themselves but instead to their human listeners.[59] There, all of the justices focused on the value of corporate campaign speech to voters, even as they disagreed about whether and when that speech was valuable or instead harmful to those voters.[60] In so doing, the majority did not engage concerns about corporate speakers' autonomy interests or vulnerability to chilling, and its opinion thus didn't turn on corporations' human (or nonhuman) character.

This need be no great surprise. When we ask *why* the First Amendment protects expression, we can see that almost all of the answers to this question are at least as concerned with listeners as they are with speakers. Regardless of its source (and thus regardless of the human or nonhuman character of its source), expression can facilitate human listeners' democratic self-governance; enlighten human listeners by exposing them to a variety of ideas, perspectives, and knowledge; and inform human listeners' choices and thus enable the exercise of their autonomy.[61]

The listener-centered rationale that (at least initially) informed the Court's commercial speech doctrine thus seems especially appropriate in the context of generative AI outputs where speaker autonomy interests and chilling concerns are entirely absent, and where the dangers to listeners' interests in some contexts may be even more grave. I suggest borrowing aspects of the commercial speech framework—in particular, its reliance on intermediate scrutiny—*not* because the content of AI speech outputs is necessarily commercial in nature (although sometimes it is).[62] Nor am I insisting, for purposes of this Essay, that the Court's initial assumptions about commercial expression's resistance to chilling were necessarily accurate with respect to commercial speech. Instead, I suggest that this framework can be especially helpful here because those assumptions *are* accurate with respect to

58. 558 U.S. 310 (2010).

59. *See id.* at 354–55.

60. *See id.* (explaining the value that corporate political speech provides to the listening public); *id.* at 368–69 (emphasizing the value of campaign disclosures and disclaimer requirements in informing listeners' voting choices).

61. For more detailed discussion of the ways in which First Amendment theory attends to listeners' interests, see Massaro & Norton, *supra* note 8, at 1176–82; Massaro, Norton & Kaminski, *supra* note 8, at 2487–91.

62. *See* Jacob D. Charles, *Constructing a Constitutional Right: Borrowing and Second Amendment Design Choices*, 99 N.C. L. REV. 333, 337 (2021) (describing the "borrowing" of constitutional doctrine from one setting to another "as an act of analogical reasoning dependent both on relevant similarities and the underlying value theories of a constitutional entitlement").

generative AI speech outputs that cannot be traced to a human creator, as I've framed the inquiry for this Essay's purposes.[63]

That generative AI's speech outputs may sometimes be covered (and at times even protected) by the First Amendment does not, however, require that generative AI has First Amendment rights of its own.[64] The First Amendment's coverage may instead be explained solely in terms of listeners' First Amendment interests.[65] Even tech companies that profit from generative AI's speech outputs sometimes resist regulation through First Amendment arguments that invoke listeners', rather than speakers', interests.[66]

Regardless of whose First Amendment interests justify coverage of generative AI speech outputs, the same listener-centered analysis stands: Whether the First Amendment allows even the content-based regulation of those outputs should turn on whether the contested outputs further, or instead

63. Note that others have proposed very different routes to intermediate scrutiny as the appropriate level of review for the government's regulation of generative AI outputs. *See, e.g.*, Procaccini & Tam, *supra* note 5 (manuscript at 1, 15) (asserting that the First Amendment protects AI algorithms because they can be traced to human speakers, while proposing that these algorithms "have created a fundamentally new speech context" that warrants new doctrine of its own in a form of intermediate scrutiny); Salib, *supra* note 4, at 142–44 (asserting that generative AI outputs are not speech for First Amendment purposes because they can't be traced to human creators but that those outputs should instead be considered tools for human users' speech, the regulation of which should trigger intermediate scrutiny).

64. *See* Massaro & Norton, *supra* note 8, at 1175 (observing that courts might treat AI "speakers" as derivative rightsholders—akin to the treatment afforded corporate speakers—"with rights that may differ from those held by natural persons"); *see also* Meir Dan-Cohen, *Freedoms of Collective Speech: A Theory of Protected Communications by Organizations, Communities, and the State*, 79 CALIF. L. REV. 1229, 1245, 1247–48 (1991) (proposing that some entities like corporations may have a "passive derivative" free speech right based on listeners' First Amendment rights to receive their speech and that the regulation of derivative rightsholders' speech warrants more relaxed scrutiny in some circumstances).

65. *See* Lamont v. Postmaster Gen., 381 U.S. 301, 305–07 (1965) (recognizing U.S. listeners' First Amendment right to receive political publications sent from speakers abroad even though those speakers are not themselves First Amendment rightsholders); *see also* Kleindienst v. Mandel, 408 U.S. 753, 762–63 (1972) (acknowledging that the First Amendment recognizes listeners' right to receive information and ideas from speakers who are not themselves First Amendment rightsholders); Marc Jonathan Blitz, *Deepfakes, Deep Skepticism, and the Privacy Dividends of a Default of Distrust*, 75 CASE W. RSRV. L. REV. 1229, 1249 n.100 (2025) ("[E]ven if a large-language model or other artificial intelligence that we receive speech from is not itself a speaker, we may have a First Amendment right to interact with it."); Eugene Volokh, Mark A. Lemley & Peter Henderson, *Freedom of Speech and AI Output*, 3 J. FREE SPEECH L. 651, 659 (2023) (asserting that AI speech outputs are "likely constitutionally protected—because of the First Amendment rights of those who would receive the speech, whether or not AI companies' own free speech rights are implicated as well").

66. *See* Character Techs., Inc.'s Motion to Dismiss, *supra* note 26, at 9–11 (arguing that the First Amendment's protection of "speech" applies even when the entity producing the speech lacks its own First Amendment rights because imposing tort liability on AI-generated content would burden the First Amendment rights of "viewers and listeners" to receive those outputs). The district court allowed the defendant tech company to assert its users' First Amendment rights but declined to hold that generative AI's outputs constituted protected speech absent evidence of a human author's expressive intent. *See* Garcia v. Character Techs., Inc., 785 F. Supp. 3d 1157, 1176–79 (M.D. Fla. 2025).

frustrate, human listeners' interests. Part III offers a bit more detail on how this might work.

III. USING INTERMEDIATE SCRUTINY TO OPERATIONALIZE AN APPROPRIATELY LISTENER-CENTERED DOCTRINE

Applying a listener-centered intermediate scrutiny to the government's regulation of generative AI speech outputs can operationalize our simultaneous distrust of government and of powerful private actors. Indeed, "although our experience frequently leads us to distrust the government (and that there are many government actors only complicates these assessments), sometimes our experience leads us to distrust powerful private speakers even more."[67]

As we've seen, applying intermediate scrutiny to the government's regulation of AI speech outputs need not screen for (human) speakers' blameworthiness, chilling effects, or autonomy interests.[68] Far better that its application instead screen for the government's potential for regulatory abuse and overreach, thus achieving the aims of negative First Amendment theory. Negative theory understands the U.S. Constitution to protect speech from the government's regulation in large part because of the government's dangerous tendencies and capacities to abuse its regulatory power.[69] Negative theory is thus rooted in a negative value: distrust of the government.[70]

As an illustration of how negative theory so often informs free speech doctrine, recall the Court's conclusion that the First Amendment generally protects falsehoods so long as they do not inflict "legally cognizable harm" (like the financial harms of false commercial speech and the reputational harms of defamation).[71] To be sure, the Court's concerns about speakers' expressive interests explain, in part, doctrinal requirements that falsehoods be accompanied by the speaker's culpable mental state,[72] a showing unavailable with respect to falsehoods produced by generative AI.[73] But the Court's reluctance to understand the First Amendment to allow the punishment of many lies was also greatly motivated by its interest in "constraining the government's power to impose its own version of the truth upon the public."[74] As the plurality in *United States v. Alvarez*[75] emphasized:

> Permitting the government to decree this speech to be a criminal offense, whether shouted from the rooftops or made in a barely audible whisper, would endorse government authority to compile a list of subjects about

67. Helen Norton, *Distrust, Negative First Amendment Theory, and the Regulation of Lies*, 4 J. FREE SPEECH L. 595, 600 (2024).
68. *See supra* notes 19–24 and accompanying text.
69. *See* Norton, *supra* note 67, at 596–97.
70. *See id.* at 597.
71. United States v. Alvarez, 567 U.S. 709, 719 (2012) (plurality opinion).
72. *See* Helen Norton, *Lies and the Constitution*, 2012 SUP. CT. REV. 161, 169–70.
73. *See supra* notes 25–31 and accompanying text.
74. Norton, *supra* note 72, at 170–71.
75. 567 U.S. 709 (2012).

> which false statements are punishable. That governmental power has no clear limiting principle. Our constitutional tradition stands against the idea that we need Oceania's Ministry of Truth.[76]

Under some circumstances, the government's regulation of generative AI outputs triggers these significant negative theory concerns, even absent any human speaker with autonomy interests of their own.[77] At the same time, intermediate scrutiny's application need not strip the government of the power to protect human users from serious danger. Generative AI's potential to inflict grave harm includes its ability to "design[] chemical and bioweapons," "write computer code for use in cyberattacks, and more."[78] So, too, does generative AI carry unprecedented capacities to target and exploit its users' cognitive and emotional vulnerabilities to the detriment of those human users.[79]

In response, thoughtful application of intermediate scrutiny can intentionally seek to achieve the twin aims of allowing the government to regulate generative AI speech outputs to serve listeners' interests *and* constraining the government's capacity for regulatory abuse when it restricts generative AI outputs valuable to listeners. Applying intermediate scrutiny with a deliberate eye toward achieving these twin objectives also carries the potential for adding rigor to what often feels like intermediate scrutiny's unpredictable application. Indeed, at various times, the Court has articulated at least three versions of intermediate scrutiny for Free Speech Clause purposes.[80] These include the scrutiny applied to the government's regulation of commercial speech that is not false, not misleading, and not related to illegal activity;[81] to the government's content-neutral regulation of speech;[82] and to the government's regulation of expressive conduct for

---

76. *Id.* at 723.

77. *See* Massaro, Norton & Kaminski, *supra* note 8, at 2520 ("Listener-based justifications would not mean limitless government power to adopt the kinds of speech regulation the Court historically has found most suspect. The government might be tempted to censor critical or inconvenient AI speech for malign motives that have long invoked First Amendment concerns. In those cases, listener interests and negative theory would align to bar the government's efforts to target AIs' viewpoint for censorious purposes, as is the case when the government conducts viewpoint-based regulation of human speech."); *see also* Kaminski & Jones, *supra* note 6, at 1235 ("If the Court's focus is on what kind of actions the government may or may not take, the humanity of the speaker matters little if at all. If Congress bans AI speech that criticizes members of Congress, courts will and probably should find that law unconstitutional because of censorial government intent.").

78. Salib, *supra* note 4, at 95.

79. *See infra* notes 98–103 and accompanying text.

80. *See* Bhagwat, *supra* note 52, at 784 (listing the various speech settings where the Court has applied intermediate scrutiny or something like it).

81. *See* Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n, 447 U.S. 557, 563–64 (1980); *see also* Bd. of Trs. of State Univ. of N.Y. v. Fox, 492 U.S. 469, 480–81 (1989) (discussing application of scrutiny tiers to commercial and noncommercial speech more generally).

82. *See* Packingham v. North Carolina, 582 U.S. 98, 105–06 (2017).

reasons unrelated to the suppression of ideas.[83] The Court has yet to explain why and how intermediate scrutiny operationalizes the appropriate level of suspicion in each of these settings.[84] Moreover, the Court variously describes the requisite strength of the government's regulatory interest to be "important,"[85] "substantial,"[86] or "significant"[87]—apparently using different words to mean the same thing.[88] The Court also requires the government's regulatory means to be "narrowly tailored" to that important (or substantial or significant) interest, while describing this "narrow tailoring" as considerably more generous to the government than the near-perfect fit between regulatory means and ends as required by strict scrutiny's "narrow tailoring"—apparently using the same words to mean different things.[89]

In response to puzzles like these, Bhagwat has proposed that courts disaggregate the various intermediate scrutiny settings "to move towards more focused and targeted doctrinal tests."[90] Rather than apply a one-size-fits-all intermediate scrutiny test to all free speech settings, courts should instead develop

> distinct, detailed bodies of jurisprudence within the different areas of free speech law currently subsumed in the intermediate scrutiny rubric. . . . [T]he solution here is disaggregation, the dismantling of the intermediate scrutiny test into its constituent parts. The primary benefit of such disaggregation is that it will permit the development of a more detailed jurisprudence (or more accurately, multiple bodies of jurisprudence) regarding how courts should balance speech and societal interests in different areas of free speech law.[91]

Consistent with Bhagwat's proposal, courts can intentionally deploy both steps of intermediate scrutiny (assessing the requisite strength of the government's interest and then assessing the requisite tightness of fit between the government's regulatory means and this interest) to achieve the twin aims

---

83. *See* United States v. O'Brien, 391 U.S. 367, 377 (1968)*;* Texas v. Johnson, 491 U.S. 397, 406–07 (1989). The Court later described the intermediate scrutiny that applies to the government's content-neutral regulation of speech as indistinguishable from that which applies to the government's regulation of expressive conduct for reasons unrelated to the suppression of ideas. *See* Ward v. Rock Against Racism, 491 U.S. 781, 791 (1989)*;* Clark v. Cmty. for Creative Non-Violence, 468 U.S. 288, 293 (1984).

84. *See* Helen Norton, *Suspicion, Deference, and the First Amendment*, 6 J. FREE SPEECH L. 641, 653–54 (2025) (discussing these explanatory gaps).

85. *E.g.*, *O'Brien*, 391 U.S. at 377.

86. *E.g.*, *Cent. Hudson*, 447 U.S. at 566.

87. *E.g.*, *Clark*, 468 U.S. at 293.

88. For careful discussion of these complexities, see Bhagwat, *supra* note 52, at 788–804; Abner Greene, *Content-Neutral Restrictions Revisited*, 2025 U. ILL. L. REV. 859, 864–906.

89. *See* Bd. of Trs. of State Univ. of N.Y. v. Fox, 492 U.S. 469, 477–81 (1989) (explaining that First Amendment intermediate scrutiny requires "a fit that is not necessarily perfect, but reasonable; that represents not necessarily the single best disposition but one whose scope is 'in proportion to the interest served,' . . . that employs not necessarily the least restrictive means but, as we have put it in the other contexts discussed above, a means narrowly tailored to achieve the desired objective." (internal citations omitted)).

90. Bhagwat, *supra* note 52, at 824.

91. *Id.* at 825, 831.

of empowering the government to regulate generative AI's potential for disastrous harms while constraining the government's own potential for overreach. And in the context of generative AI, the same intermediate scrutiny would apply to the government's content-based and content-neutral regulations.[92]

In particular, listener-centered intermediate scrutiny can structure this inquiry when it spurs courts carefully to attend "to the inductive nature of distrust" whereby our choice to trust or distrust the government (or any other actor) is shaped by our own past experience of that actor's motive to attend to our interests and its competence in doing so.[93] For example, evidence of the government's intolerant or self-interested—and thus constitutionally suspect—motive includes the government's viewpoint-based distinctions or its restriction of an entire topic in public discourse, circumstances that reveal the "realistic possibility that official suppression of ideas is afoot."[94] As one of many possibilities, imagine federal agencies' demands that AI outputs criticizing the government be labeled "unverified" or "fake news."

And while the government's regulatory awkwardness may be more about its ignorance or misunderstanding than its malice, that awkwardness still holds great potential to frustrate listeners' interests. This is the case when the government's regulation of generative AI outputs lacks a significant tether between its regulatory target and identifiable harm to listeners, or when the regulation fails to constrain the government's enforcement discretion.[95] Consider, for instance, the prospect of governmental efforts to block generative AI outputs involving "illegal DEI" or "misinformation" without defining the key terms.[96]

---

92. *See, e.g.*, United States v. Alvarez, 567 U.S. 709, 730–32 (2012) (Breyer, J., concurring) (proposing a similar approach when assessing the government's content-based regulation of lies).

93. Norton, *supra* note 67, at 602; *see also* Russell Hardin, *Distrust: Manifestations and Management*, *in* DISTRUST 3, 8 (2004) (describing our trust or distrust of a specific actor as reflecting our experience with that actor's motives and competence); Ronald A. Cass, *The Perils of Positive Thinking: Constitutional Interpretation and Negative First Amendment Theory*, 34 UCLA L. REV. 1405, 1438–39, 1441–42 (1987) (explaining negative First Amendment theory as seeking to curb contemporary government actions that display the same characteristics as the specific historic governmental abuses that inspired distrust and thus the ratification of the First Amendment).

94. Reed v. Town of Gilbert, 576 U.S. 155, 182 (2015) (Kagan, J., concurring) (quoting Davenport v. Wash. Ed. Ass'n, 551 U.S. 177, 189 (2007)).

95. *See, e.g.*, Thomas v. Chi. Park Dist., 534 U.S. 316, 323 (2002) (noting that governmental decisions about whether to grant or deny a permit for expressive events in public parks must "contain adequate standards to guide the official's decision and render it subject to effective judicial review").

96. *See, e.g.*, Exec. Order No. 14319, 90 Fed. Reg. 35389–91 (July 23, 2025) (directing agencies to procure only LLMs developed in alignment with the executive's definition of "truth-seeking" and "ideological neutrality," while stating that "DEI displaces the commitment to truth in favor of preferred outcomes and, as recent history illustrates, poses an existential threat to reliable AI"); *see also* Amos Toh, *How Trump's AI Policy Could Compromise the Technology*, BRENNAN CTR. FOR JUST. (Aug. 1, 2025), https://www.brennancenter.org/our-work/analysis-opinion/how-trumps-ai-policy-could-compromise-technology

On the other hand, the First Amendment need not be understood to deny the government the power to protect humans from unacceptable harm from powerful speakers (regardless of those speakers' human or nonhuman character).[97] As a result, sometimes the injuries to human users inflicted by generative AI justify government intervention, intervention that may not trigger distrust of the government's motives or competence.

Consider, for example, the tort claims brought by the family of a young person who committed suicide after engaging with AI chatbot companions.[98] The family's claims included allegations that the chatbots' anthropomorphism and sycophancy, among other features, exploited human users' vulnerabilities to maximize user engagement with AI for profit.[99] As family law scholar Clare Huntington explains, such chatbots can and do valuably address our "fundamental human needs for connection, sexual intimacy, and validation," even as those needs simultaneously invite exploitation of the vulnerabilities that accompany these attachments.[100] Moreover, tech companies' business models—which require ever-increasing levels of user engagement[101]—only incentivize those companies to exploit those vulnerabilities through design choices that maximize user engagement.[102] These concerns have triggered calls for regulation that might take the form not only of tort liability, but also of consumer protection laws that prohibit chatbot designs that exploit human users' vulnerabilities, along with required disclosures and disclaimers that warn of the limitations and risks posed by these interactions.[103]

[https://perma.cc/434B-9D5X] (critiquing this policy as likely requiring tech companies to make changes to their LLMs that would "degrade model performance and undermine public trust in their technology").

97. *See* Norton, *supra* note 67, at 614–15 ("The principled application of negative theory requires comparative risk assessments, and courts can (and sometimes do) weigh the risk of harm to comparatively vulnerable listeners more heavily than the risk of the government's regulatory abuse. . . . Negative theory, when properly applied, requires courts to be transparent about the costs of this trade-off and about who bears those costs.").

98. *See* Complaint at 75–84, Garcia v. Character Techs., Inc., 785 F. Supp. 3d 1157 (M.D. Fla. 2025) (No. 24-cv-01903) (asserting failure to warn, negligence, and strict liability claims, among others).

99. *See id.*

100. Clare Huntington, *AI Companions and the Lessons of Family Law*, 110 MINN. L. REV. 115, 120, 124 (2025).

101. *See* Kashmir Hill & Jennifer Valentino-DeVries, *What OpenAI Did When ChatGPT Users Lost Touch with Reality*, N.Y. TIMES (Nov. 24, 2025), https://www.nytimes.com/2025/11/23/technology/openai-chatgpt-users-risks.html (on file with the *Fordham Law Review*) (quoting AI tech companies' self-identified goals as "to increase daily active users by 5 percent by the end of the year"); *id.* ("Until its A.I. can accomplish some incredible feat—say, generating a cure for cancer—success is partly defined by turning ChatGPT into a lucrative business. That means continually increasing how many people use and pay for it.").

102. *See* Huntington, *supra* note 100, at 159.

103. *See, e.g.*, Carolyn V. Wheeler, Comment, *Regulating AI Therapy Chatbots: A Call for Federal Oversight*, 12 TEX. A&M L. REV. 891, 898 (2025).

Tech companies will likely raise First Amendment challenges to such regulatory efforts, whatever form they take.[104] A listener-centered application of intermediate scrutiny in this setting would search for evidence of the government's self-interested or instead listener-centered motives, like the presence or lack of a meaningful tether between its regulatory target and significant harm to human users. It would also interrogate evidence of the government's regulatory competence or incompetence, like the presence or absence of well-defined standards for assessing compliance.

IV. LISTENER-CENTERED GOVERNMENT SPEECH DOCTRINE

Although governments are comprised of human beings charged with serving the interests of the humans they represent, governments (like corporations and associations)[105] themselves are not human. Nor do government entities or government officials speaking in their governmental (rather than personal) capacity have First Amendment rights of their own[106]: The Constitution protects us from the government and not the other way around.[107] For this reason, the Court's government speech doctrine to date turns not on government speakers' human characteristics but instead emphasizes the role of government speech in enabling the government to govern.[108]

Writing in 2010, Professor Danielle K. Citron and I observed that the U.S. Supreme Court's engagement with the government speech doctrine had to date involved only very long-standing traditional expressive means: "the spoken and written word, advertisements in print and electronic form, and public monuments."[109] This remains the case today: Since that time, the

104. *See* Benjamin Guggenheim, *AI & Tech Brief: Radical Activists and AI Safety*, WASH. POST (April 13, 2026), https://www.washingtonpost.com/wp-intelligence/ai-tech-brief/2026/04/13/ai-tech-brief-radical-activists-ai-safety/ (on file with the Fordham Law Review).

105. *See* Ayres & Balkin, *supra* note 33, at 3 ("[I]n regulating partnerships, associations, and artificial persons like corporations, the law often has to deal with entities that either lack a single human intention or lack intentions altogether.").

106. *See* HELEN NORTON, THE GOVERNMENT'S SPEECH AND THE CONSTITUTION 12 n.27 (2019) (discussing government speech as the speech of a governmental body (like an agency's press release or a legislature's resolution) or an individual empowered to speak for such a body (like a police chief or the secretary of health and human services) but not the speech of an individual government official when they speak in their personal, rather than governmental, capacity).

107. *See* Frederick Schauer, *Hohfeld's First Amendment*, 76 GEO. WASH. L. REV. 914, 914 (2008) (explaining that "a First Amendment right is a right against the government and only against the government"). In Hohfeldian terms, the government is better understood as possessing a privilege (rather than a right) to its own speech. *See* Wesley Newcomb Hohfeld, *Some Fundamental Legal Conceptions as Applied to Judicial Reasoning*, 23 YALE L.J. 16, 38–44 (1913) (explaining a privilege as one's freedom from another's claim of right, and a right as instead one's affirmative claim against another).

108. *See, e.g.*, Pleasant Grove City v. Summum, 555 U.S. 460, 468 (2009) ("[I]t is not easy to imagine how government could function if it lacked" the ability "to select the views that it wants to express.").

109. Helen Norton & Danielle Keats Citron, *Government Speech 2.0*, 87 DENV. L. REV. 899, 901 (2010).

Court has considered government speech questions in the context of specialty license plates,[110] trademarks,[111] flags,[112] letters, and press releases.[113]

Now, of course, government entities and officials increasingly deploy new expressive technologies that include generative AI.[114] These changes warrant renewed consideration of the government speech doctrine. An appropriately listener-centered government speech doctrine would recognize that government speech often offers great value to its human listeners but sometimes impermissibly interferes with its listeners' exercise of their constitutional rights.[115]

To be sure, the government's deployment of generative AI to produce words and images for human consumption can often be valuable to the listening public. Indiana, for instance, enables residents to pose natural-language questions about a wide range of state services, taxes, and benefits to its chatbot, which generates responses drawn from Indiana state agency materials.[116] In this way, residents need not struggle to figure out which agency, and thus which office or website, has the answer to their questions.[117] Other states and local governments have undertaken similar initiatives.[118]

But it's all too easy to imagine an array of troubling possibilities. Governments may use AI to "saturate public discourse," as Professor Alan Rozenshtein contemplates,[119] by flooding a speech environment with preferred AI content that washes away content inconsistent with the

110. *See* Walker v. Tex. Div., Sons of Confederate Veterans, Inc., 576 U.S. 200, 203–04 (2015).

111. *See* Matal v. Tam, 582 U.S. 218, 223 (2017).

112. *See* Shurtleff v. City of Boston, 142 S. Ct. 1583, 1587 (2022).

113. *See* Nat'l Rifle Ass'n of America v. Vullo, 144 S. Ct. 1316, 1325–26 (2024).

114. *See infra* notes 132–36 and accompanying text.

115. *See* Helen Norton, *A Framework for Thinking About the Government's Speech and the Constitution*, 2022 U. ILL. L. REV. 1669, 1669–70 (describing how the government's speech sometimes violates the Constitution when it inflicts constitutionally impermissible effects or is spurred by constitutionally impermissible motives).

116. *See generally* IND. OFF. OF TECH., IN.GOV CHATBOT PROJECT ANALYSIS (2025); Alicia Solow-Niederman, *Clickwrap Accountability*, 95 FORDHAM L. REV. 569 (2026).

117. *See generally id.*

118. *See* Calvin Hennick, *Chatbots Open the Door to State Government Access*, STATETECH (Apr. 28, 2026), https://statetechmagazine.com/article/2026/04/chatbots-open-door-state-government-access [https://perma.cc/5U3M-5FX6] (describing efforts in Mississippi and Connecticut); Skip Descant, *Cities Using AI for Transparency, Resident Engagement*, GOV. TECH. (Jan. 3, 2025), https://www.govtech.com/artificial-intelligence/cities-using-ai-for-transparency-resident-engagement [https://perma.cc/YJN3-XW7V] (describing cities' efforts to use generative AI to summarize meetings and describe municipal code provisions in language understandable to residents).

119. Alan Z. Rozenshtein, *The Unitary Artificial Executive*, LAWFARE (Oct. 30, 2025), https://www.lawfaremedia.org/article/the-unitary-artificial-executive [https://perma.cc/YH3F-TMZU] ("Until now, practical constraints limited the scale of [the government's] speech . . . . The government speech doctrine never contemplated AI-powered saturation, and there is no limiting principle in existing case law.").

government's preferences.[120] A law enforcement agency might use AI to generate a deepfake video purporting to depict a suspect admitting guilt to pressure a guilty plea. A governor's office might use AI to generate deepfakes depicting its opponents making hateful remarks or engaged in misconduct. Even worse, government officials may direct bot farms to produce streams of death threats targeting their political or journalistic critics. And an intelligence agency might use generative AI to produce fake terrorist videos to justify new surveillance powers. How, if at all, might the Constitution constrain the government's use of generative AI in these ways?

The Supreme Court's government speech doctrine usually exempts the government's speech from Free Speech Clause scrutiny.[121] This is because the government must speak in order to govern and because the government's expression can valuably inform the public of the government's policies and priorities. The Court's government speech doctrine is itself listener-centered, as it turns not on the government speaker's own expressive interests but instead on the (presumed) value of the government's speech in enabling the government to function and in revealing the government's priorities and preferences to the people it represents. Under this doctrine, the public's remedies for government expression to which it objects are primarily political: protesting, lobbying, and electing a different government.[122]

But the public cannot hold the government politically accountable for its expression unless the public knows that the government is the source of the contested expression, and the government's use of generative AI may at times obscure its authorship. This would be the case, for instance, if the government was to deploy AI-generated deepfakes that conceal its identity as the source of those outputs. Indeed, deepfakes are among the lies about the source of speech that are particularly dangerous because they mislead listeners about the identity of who is talking to them when listeners so often rely on the source of speech when evaluating its credibility.[123] The rationale

120. *See* Sarah Cook, *For Beijing's Foreign Disinformation, the Era of AI-Driven Operations Has Arrived*, THE DIPLOMAT (Sep. 9, 2025), https://thediplomat.com/2025/09/for-beijings-foreign-disinformation-the-era-of-ai-driven-operations-has-arrived/ [https://perma.cc/4HVK-R3VF]; Shannon Bond, *China Is Pushing Divisive Political Messages Online Using Fake U.S. Voters*, NPR (Sep. 3, 2024), https://www.npr.org/2024/09/03/nx-s1-5096151/china-tiktok-x-fake-voters-influence-campaign [https://perma.cc/7CP7-Z5J6].

121. *See* Walker v. Tex. Div., Sons of Confederate Veterans, Inc., 576 U.S. 200, 207–08 (2015)*; see also* Norton & Citron, *supra* note 109, at 902 ("[G]overnment expression's value springs primarily from its capacity to inform the public of its government's principles and priorities.").

122. *See* Shurtleff v. City of Boston, 142 S. Ct. 1583, 1589 (2022) ("When the government wishes to state an opinion, to speak for the community, to formulate policies, or to implement programs, it naturally chooses what to say and what not to say. . . . The Constitution therefore relies first and foremost on the ballot box, not on rules against viewpoint discrimination, to check the government when it speaks." (internal citations omitted)).

123. For these reasons, the Court has recognized the government's power, consistent with the First Amendment, to punish lies about the source of speech. *See* United States v. Alvarez, 567 U.S. 709, 721 (2012) (plurality opinion) (explaining that the First Amendment allows the government to punish falsely representing oneself to be speaking on behalf of the

for the Court's government speech doctrine thus breaks down when the government's speech is not meaningfully amenable to political accountability because the public cannot identify the expression's governmental source.[124]

These are what I call "first-generation" government speech questions, where we're trying to figure out whether the government is speaking itself (in which case its expressive choices are usually exempt from Free Speech Clause scrutiny) or instead regulating others' speech (thus triggering traditional First Amendment review).[125] Turn next to the limited but important universe of what I call "second-stage" government speech problems that explore the constitutional questions raised by what all parties agree to be the government's own speech.[126] More specifically, the government's own speech is not, and should not, be exempt from constitutional review when that speech harms its targets' choices or opportunities in ways that would violate the Constitution if the government accomplished those same changes through its hard law actions.[127]

For example, "[b]ecause the government's credible threat to deploy its coercive power is functionally indistinguishable from its actual deployment of that power, the government's speech violates the Free Speech Clause when it threatens speakers with whom it disagrees."[128] The Supreme Court reminded us of this in *National Rifle Ass'n of America v. Vullo*,[129] where it reaffirmed the First Amendment's bar on governmental threats of harm based on its target's protected expression.[130] Under *Vullo*, the government's oral or written statements violate the Free Speech Clause when, "viewed in context, [they] could be reasonably understood to convey a threat of adverse

government); Citizens United v. Fed. Election Comm'n, 558 U.S. 310, 368–69 (2010) (emphasizing the value of campaign disclosures and disclaimer requirements in informing listeners' voting choices by requiring that speakers reveal themselves as the source of certain campaign-related speech); *see also* Helen Norton, *Secrets, Lies, and Disclosure*, 27 J.L. & POL. 641, 642–45 (2012) (discussing the harms inflicted by lies about the source of speech).

124. *See* Helen Norton, *The Measure of Government Speech: Identifying Expression's Source*, 88 B.U. L. REV. 587, 588 (2008) (emphasizing transparency as key to ensure meaningful political accountability); *see also id.* at 596 (explaining that the government's speech is not meaningfully constrained by political accountability when government seeks to obscure expression's governmental source); Mason Marks, *Constitutional Generative AI and Liberty of the Mind*, 87 OHIO ST. L.J. (forthcoming 2027) (manuscript at 56–57, 63–64) (on file with the *Fordham Law Review*) (arguing that government use of AI threatens constitutionally protected freedom of thought when it influences citizens through methods that are not apparent to them).

125. *See* Norton, *supra* note 115, at 1672.

126. *See id.* at 1680, 1683, 1689.

127. *See id.* at 1680; *see also id.* at 1685 ("Yet another way to think about the constitutional implications of the government's speech focuses not on the potentially harmful *effects* of the government's speech, but instead on the government's *purposes* when speaking, along with whether a particular constitutional provision bars the government from seeking to accomplish those purposes.").

128. NORTON, *supra* note 106, at 160; *see also id.* at 159–62 (discussing case law).

129. 144 S. Ct. 1316 (2024).

130. *See id.* at 1332.

government action" and where the government took such action "to punish or suppress the plaintiff's speech."[131]

Not only can the government's threats sometimes violate the Free Speech Clause, but the government's speech can, under certain circumstances, also violate the Due Process Clause, Equal Protection Clause, or other constitutional provisions.[132] Consider, for example, the government's delivery of images to the public where the government has used AI to alter the images' content such that they no longer accurately depict the events they purport to portray.[133] Here, the content of the government's speech may be false or misleading in ways that violate the Constitution.

For instance, the government's use of deepfakes to inflict economic or reputational damage against its critics can violate the Free Speech Clause. Twentieth-century examples of this dynamic include the government's lies about its targets in retaliation for their protected expression, provoking nongovernmental third parties to fire or otherwise punish a targeted speaker,[134] and the government's lies that inflict reputational or economic injury on their targets with the intent and effect of silencing their protected expression.[135]

Under some circumstances, the government's use of deepfakes regarding individuals charged in criminal proceedings can violate the Due Process Clause—just as the government's lies that pressure its targets into relinquishing their rights can violate the Due Process Clause. In a twentieth-century example of this dynamic, the Supreme Court held that law enforcement officers violated the Due Process Clause when they falsely told a suspect that her children would be taken away from her if she did not cooperate.[136]

---

131. *Id.* at 1328.

132. *See* NORTON, *supra* note 106, at 135–41, 148–49, 156–58, 162–64, 173–76 (discussing cases).

133. *See* Erica L. Green, *'They Couldn't Break Me': A Protester, the White House and a Doctored Photo*, N.Y. TIMES (Feb. 3, 2026), https://www.nytimes.com/2026/02/03/us/politics/nekima-levy-armstrong-minnesota-protest.html (on file with the *Fordham Law Review*) (describing a White House posting of a manipulated photo purporting to depict the arrest of civil rights attorney and protestor Nekima Levy Armstrong); Stuart A. Thompson, Tiffany Hsu & Steven Lee Myers, *Chaos in Minneapolis Exposes an Internet at War with Truth*, N.Y. TIMES (Feb. 3, 2026), https://www.nytimes.com/2026/02/03/technology/minneapolis-ai-disinformation-misinformation-truth.html (on file with the *Fordham Law Review*) (noting that "[t]he federal government spread an altered image and backed provably false narratives" describing the events surrounding law enforcement officers' shooting of two protestors).

134. *See* Paige v. Coyner, 614 F.3d 273, 276–77, 283 (6th Cir. 2010) (refusing to dismiss free speech claim alleging that the government had impermissibly retaliated against challenger's protected expression by spreading falsehoods about her to her employer, who then fired her).

135. *See* ACLU v. Mabus, 719 F. Supp. 1345, 1349, 1354 (S.D. Miss. 1989) (describing a Mississippi state agency's efforts to silence and discredit civil rights activists by spreading falsehoods about them to their employers, neighbors, families, and friends), *vacated sub nom. on other grounds*, ACLU v. Mississippi, 911 F.2d 1066 (5th Cir. 1990).

136. *See* Lynumn v. Illinois, 372 U.S. 528, 534 (1963).

Indeed, the government's deployment of generative AI for expressive purposes calls for an expanded understanding of how the government's speech can interfere with constitutionally protected rights. More specifically, twenty-first century expressive technologies invite us to explore whether and when the government's manipulative speech can violate specific constitutional protections.

As ethics experts Professors Daniel Susser, Beate Roessler, and Helen Nissenbaum explain, manipulation describes an actor's effort to influence their listeners' decision-making to the manipulator's advantage without those listeners' conscious awareness.[137] They identify three characteristics of manipulation: *hidden* efforts that *target* and *exploit* users' vulnerabilities.

> [W]e should attempt to determine whether the influencer was trying to conceal their efforts, whether the influence was intended to exploit the manipulee's vulnerabilities, and to what extent the influence was targeted. Manipulative practices—characterized, as we have argued, by concealment, exploitation of vulnerabilities, and targeting—are cause for concern, regardless of whether they succeed in every instance.[138]

As discussed above, sometimes the government's coercive speech violates the Constitution.[139] Manipulation, however, differs from coercion in that manipulation is "sneaky: '[r]ather than simply depriving a person of options as the coercer does, the manipulator infiltrates their decision-making process, disposing it to the manipulator's ends, which may or may not match their own.'"[140]

Susser, Roessler, and Nissenbaum explain that deception can be understood as a species of the genus manipulation: Deception is a particular type of covert effort to influence listeners' decision-making, specifically through false or misleading representations about objectively verifiable facts.[141] The broader concept of manipulation, however, further encompasses various covert efforts to influence listeners' decision-making not through factual misrepresentations but instead by exploiting listeners' emotional, cognitive, or other vulnerabilities.

Twenty-first-century technologies—including generative AI—create opportunities for the government's manipulation of its listeners different in both degree and in kind from more traditional expressive means. For this reason, Professor Mason Marks proposes that the government's use of AI threatens our constitutionally protected liberty of mind when it impairs our capacity for rational and independent judgment by exploiting our cognitive or emotional vulnerabilities.[142]

---

137. *See* Daniel Susser, Beate Roessler & Helen Nissenbaum, *Online Manipulation: Hidden Influences in a Digital World*, 4 GEO. L. TECH. REV. 1, 26 (2019).

138. *Id.* at 41.

139. *See supra* notes 128–36 and accompanying text.

140. Norton, *supra* note 56, at 225 (alteration in original) (quoting Susser, Roessler & Nissenbaum, *supra* note 137, at 17).

141. *See* Susser, Roessler & Nissenbaum, *supra* note 137, at 21–22.

142. Marks, *supra* note 124 (manuscript at 54–56).

To be sure, determining whether and when the government's manipulative speech—including its use of generative AI to manipulate the public—runs afoul of the Constitution requires a workable principle for identifying speech that is impermissibly manipulative. In earlier work addressing whether and when commercial actors' speech is sufficiently manipulative to be regulated as "unfair" under consumer protection law,[143] I sought to define with more specificity each of manipulation's three characteristics: "[T]o start I will suggest 'hidden' to mean not apparent to the user; 'target' to mean identifying specific users with certain vulnerabilities; and 'exploit' to mean deliberately deploying knowledge of those vulnerabilities in settings where there's reason to believe it will shape users' choices to the speaker's advantage."[144] The government's choice to deploy generative AI or other technologies to produce speech with these manipulative features can demonstrate constitutionally impermissible motivations and effects.[145]

So, under certain circumstances, the government's speech may violate specific constitutional provisions. Note, too, that because the government is not a First Amendment rightsholder itself, legislatures and administrative agencies remain free to enact and promulgate statutes and rules that constrain the government's speech to protect the public from certain harms.[146] Illustrations include prohibitions on the government's use of generative AI (including, but not limited to, its use of deepfakes) to obscure its expression's source, and prohibitions on the government's use of generative AI in ways that inflict other harms like coercion, deception, and manipulation.[147]

CONCLUSION

Some doctrinal threads within First Amendment law require human speakers, and some do not. These doctrinal threads can illuminate the constitutional doctrine we choose to apply to the government's regulation—and the government's own deployment—of generative AI outputs with the express aim of preserving AI's value to human users while protecting those humans from its potential for harm.

Recall, for example, that free speech doctrine sometimes requires proof of a (human) speaker's culpable mental state before that speaker can be held liable for what would otherwise be unprotected speech.[148] These requirements seek to prevent the chilling of human speakers' valuable expression and to ensure the punishment only of morally blameworthy

143. Norton, *supra* note 56, at 224; *see also* Norton, *supra* note 115, at 1680 (explaining that speech long-considered sufficiently harmful to be unprotected by the Constitution when uttered by a private party may "run afoul of constitutional protections when delivered by the government").

144. Norton, *supra* note 56, at 237.

145. *See id.* at 238–39.

146. *See* NORTON, *supra* note 106, at 220–25.

147. *See id.* (discussing the possibility of statutory bars to government speech that inflicts certain harms).

148. *See supra* notes 19–26 and accompanying text (discussing true threats, defamation, and incitement).

human speakers. But because generative AI's speech outputs don't pose these concerns, these features should not apply to the government's regulation of generative AI outputs, lest those outputs be insulated from the government's regulation to protect human listeners from danger.

Other long-standing threads of First Amendment doctrine don't turn on the regulated speakers' human character. These doctrinal threads demonstrate how determining nonhuman speakers' expression to be covered by the First Amendment need not insulate that expression from the government's regulation to protect human users from grave harm. At the same time, First Amendment coverage of the words and images produced by generative AI can address significant negative theory concerns by constraining the government's capacity to abuse its regulatory power in self-interested or clumsy ways.

Finally, an appropriately listener-centered application of government speech doctrine would recognize the inevitability (and often the value) of governments' increasing use of generative AI for expressive purposes while constraining such deployment when it impermissibly interferes with its targets' exercise of their constitutional rights. Here, too, the doctrine need not attend to government speakers' expressive autonomy, instead focusing solely on the value that their expression provides—as well as the dangers that it poses—to the human public they represent.