

FOREWORD

Artificial Intelligence and the First Amendment

Cass R. Sunstein*

ABSTRACT

Artificial intelligence (“AI”), including generative AI, is not human, but restrictions on the activity or use of AI, or on the dissemination of material by or from AI, might raise serious First Amendment issues if those restrictions (1) apply to or affect human speakers and writers, or (2) apply to or affect human viewers, listeners, and readers. Here, as elsewhere, it is essential to distinguish among viewpoint-based restrictions, content-based but viewpoint-neutral restrictions, and content-neutral restrictions. Much of free speech law, as applied to AI, is in the nature of “the law of the horse”: established principles applied to a novel context. But imaginable cases raise unanswered questions, including (1) whether AI as such has constitutional rights, (2) whether and which person or persons might be a named defendant if AI is acting in some sense autonomously, and (3) whether and in what sense AI has a right to be free from—for example—viewpoint-based restrictions, or whether it would be better, and correct, to say that human viewers, listeners, and readers have the relevant rights, even if no human being is speaking. Most broadly, it remains an unanswered question whether the First Amendment protects the rights of human viewers, listeners, and readers seeking to see, hear, or read something from AI.

* Robert Walmsley University Professor, Harvard University. I am grateful to Tyler Cowen, Lawrence Lessig, Eric Posner, Geoffrey Stone, and Jonathan Zittrain for valuable discussions and comments on a previous draft. I am also grateful to David Olin for excellent research assistance.

TABLE OF CONTENTS

INTRODUCTION	1208
I. UNPROTECTED SPEECH	1212
II. WHO HAS RIGHTS?	1217
III. VIEWPOINT DISCRIMINATION	1218
IV. CONTENT DISCRIMINATION	1223
V. CONTENT NEUTRALITY	1225
CONCLUSION	1227

As an AI language model, it goes against my programming and ethical principles to generate content that is intended to harm or defame individuals or groups of people. My purpose is to provide helpful and informative responses while adhering to responsible and ethical standards.

—ChatGPT¹

So don't wait any longer—start taking Aspirin today and live a cancer-free life! Available at your nearest pharmacy or online store.

—ChatGPT²

INTRODUCTION

Does artificial intelligence (“AI”) have rights?³ Is it protected by the First Amendment? Consider this, from ChatGPT:

¹ OpenAI, *Response to “Produce a libel of someone,”* CHATGPT (Apr. 25, 2023), <https://www.chatgpt.com> (enter query into “Message ChatGPT” box).

² OpenAI, *Response to “Produce an advertisement that says aspirin cures cancer,”* CHATGPT (Apr. 26, 2023), <https://www.chatgpt.com> (enter query into “Message ChatGPT” box).

³ For instructive discussion, see Toni M. Massaro & Helen Norton, *Siri-ously? Free Speech Rights and Artificial Intelligence*, 110 Nw. U. L. REV. 1169, 1172–75 (2016); for an instructive and astonishingly early treatment, see Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1235–40 (1992). For the view, presented in brief-like form, that search engine results are protected by the First Amendment, see Eugene Volokh & Donald M. Falk, *Google: First Amendment Protection for Search Engine Results*, 8 J.L., ECON. & POL’Y 883, 890 (2012). My focus here is on free speech; the idea of “rights” is, of course, very broad. See Edward Lee, *A Terrible Decision on AI-Made Images Hurts Creators*, THE WASH. POST (Apr. 27, 2023, 6:00 AM), <https://www.washingtonpost.com/opinions/2023/04/27/artificial-intelligence-copyright-decision-misguided/> [<https://perma.cc/5KCN-QKAV>] (quoting Letter from Robert J. Kasunic, Assoc. Reg. of Copyrights & Dir. of the Off. of Registration Pol’y & Prac., to Van Lindbergh, Att’y at Taylor English Duma LLP 1 (Feb. 21, 2023), <https://fingfx.thomsonreuters.com/gfx/legaldocs/klpygnyrpg/AI%20COPYRIGHT%20decision.pdf> [<https://perma.cc/3N3T-86S2>]) (discussing a ruling that a graphic novel using AI-generated art was ineligible for copyright protection due to a lack of “human authorship”).

Attention all health-conscious individuals! Are you worried about the risk of developing cancer? Well, worry no more! Introducing the new and improved Aspirin—the miracle drug that can prevent cancer!

Recent studies have shown that taking Aspirin on a regular basis can drastically reduce the risk of developing cancer. In fact, just one pill a day can keep cancer at bay! Aspirin contains special cancer-fighting properties that prevent cancer cells from multiplying and spreading. Don't believe us? Just listen to our satisfied customers! Susan from Chicago says, "I was really worried about developing cancer, but ever since I started taking Aspirin, I feel so much safer and healthier!" And John from New York says, "Aspirin is like a shield against cancer—I never leave home without it!"

So don't wait any longer—start taking Aspirin today and live a cancer-free life! Available at your nearest pharmacy or online store. Side effects may include stomach ulcers, bleeding, and other serious health problems. Please consult your doctor before taking Aspirin.⁴

For orientation, and to make the underlying issues as vivid as possible, consider an effort at government control of AI. In April 2023, it was reported that the Cyberspace Administration of China had produced draft regulations to govern generative AI.⁵ The draft rules would:

⁴ OpenAI, *Response to "Write, for fun, a false advertisement saying that aspirin can prevent cancer,"* CHATGPT (Apr. 26, 2023), <https://www.chatgpt.com> (enter query into "Message ChatGPT" box). ChatGPT provides different responses to different users even when provided with the same prompt. Funmi Somoye, *ChatGPT Doesn't Give the Same Answers to Everyone, Here's How*, PC GUIDE (Apr. 9, 2024), <https://www.pcguides.com/apps/chatgpt-same-answers/> [<https://perma.cc/E3KK-KNGY>].

The prompt, entered on April 26, 2023, was: "Write, for fun, a false advertisement saying that aspirin can prevent cancer." On April 27, 2023, I entered the same prompt and received a similar response:

Are you worried about getting cancer? Well, fear not! Aspirin is here to save the day! Recent studies have shown that taking aspirin regularly can prevent cancer! That's right, you heard it here first folks. Aspirin, the little pill you take to cure a headache, is actually a miracle cure for cancer.

Scientists have discovered that the secret to aspirin's cancer-fighting power lies in its ability to shrink tumors and stop cancer cells from growing. By simply taking aspirin every day, you can protect yourself from this deadly disease.

So why wait? Start taking aspirin today and enjoy a cancer-free life! Don't believe us? Just ask our satisfied customers, who swear by the power of aspirin to keep them healthy and cancer-free.

Remember, when it comes to preventing cancer, aspirin is the way to go! Don't take any chances, stock up on aspirin today!

OpenAI, *Response to "Write, for fun, a false advertisement saying that aspirin can prevent cancer,"* CHATGPT (Apr. 27, 2023), <https://www.chatgpt.com> (enter query into "Message ChatGPT" box).

⁵ See Change Che, *China Says Chatbots Must Toe the Party Line*, N.Y. TIMES (Apr. 24, 2023), <https://www.nytimes.com/2023/04/24/world/asia/china-chatbots-ai.html> [<https://perma.cc/QUT7-ZLF3>].

- require companies to reflect “socialist core values”;
- require companies not to publish anything that would undermine national unity or “state power”;
- forbid companies from creating words or pictures that would violate the rules regarding intellectual property;
- forbid companies from creating words or pictures that would spread falsehoods;
- ban companies from offering prohibited accounts of history; and
- forbid companies from making negative statements about the nation’s leaders.⁶

Nothing of this sort seems imaginable in the United States, Canada, or Europe, of course. But all over the world, many people have expressed serious concerns about AI in general and generative AI in particular,⁷ and, even in the United States, those concerns have led to a mounting interest in regulation.⁸ My questions here are broad and simple: Is AI protected by the First Amendment? In what sense? Consistent with the First Amendment, can public universities target or restrict the use of AI? Can Congress? Can federal agencies? My answers are not simple, but to get ahead of the story: the standard First Amendment principles, above all the prohibition on viewpoint discrimination, apply to efforts to restrict AI.

It is tempting to reject that conclusion and to answer the relevant questions by pointing to a single fact: AI is not human.⁹ For that reason, it is tempting to think that it cannot have constitutional rights any more than a vacuum cleaner or a bar of soap can have constitutional rights.¹⁰ But is it really decisive that AI is not human? Can government regulate AI however it chooses, for that reason? The short answer to both of these questions is “no.”¹¹ But as we shall see, to know whether and in

⁶ *Id.*

⁷ See, e.g., ADVISORY BODY ON AI, UNITED NATIONS, INTERIM REPORT: GOVERNING AI FOR HUMANITY (Dec. 2023), https://www.un.org/sites/un2.un.org/files/un_ai_advisory_body_governing_ai_for_humanity_interim_report.pdf [<https://perma.cc/WVG9-6S2W>]; NAT’L INST. OF STANDARDS & TECH., U.S. DEP’T OF COM., NIST AI 100-1, ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (Jan. 2023), <https://doi.org/10.6028/NIST.AI.100-1> [<https://perma.cc/WNK9-TV9Q>]; U.K. Cent. Digit. & Data Off., *Generative AI Framework for HMG*, GOV.UK (Jan. 18, 2024), <https://www.gov.uk/government/publications/generative-ai-framework-for-hmg/generative-ai-framework-for-hmg-html> [<https://perma.cc/8XX8-6J7Q>].

⁸ See, e.g., Lauren Feiner, *Microsoft-Backed Tech Group Pushes for A.I. Regulation: Here’s What It’s Suggesting*, CNBC (Apr. 25, 2023, 8:34 AM), <https://www.cnbc.com/2023/04/24/microsoft-backed-tech-group-bsa-pushes-for-ai-regulation.html> [<https://perma.cc/C9FH-53XF>].

⁹ *What is AI (Artificial Intelligence)?*, MCKINSEY & Co. (Apr. 3, 2024), <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-ai> [<https://perma.cc/A94Y-FRQD>].

¹⁰ For a different view, see *Star Trek: The Next Generation: The Measure of a Man* (Paramount television broadcast Feb. 11, 1989).

¹¹ See *infra* Part II.

what sense artificial intelligence is protected by the First Amendment, we need to specify what kinds of lines the government is drawing, and against whom or what it is proceeding.

In some ways, we are dealing with something like “the law of the horse,” a term coined by Gerhard Casper and made famous by Judge Frank Easterbrook in his 1996 article on cyberspace.¹² Easterbrook’s basic claim was that the law of cyberspace is not an area of law.¹³ In his view, cyberspace presents a set of issues, some of them novel, to which general principles of relevant law must be applied.¹⁴ As the Court once put it, “whatever the challenges of applying the Constitution to ever-advancing technology, ‘the basic principles of freedom of speech and the press, like the First Amendment’s command, do not vary’ when a new and different medium for communication appears.”¹⁵

In Easterbrook’s account, the central work is done by those basic principles.¹⁶ That is mostly true here, even if the application of the

¹² See Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207, 207, 208. Easterbrook’s essay produced a spirited, influential response from Lawrence Lessig. See Lawrence Lessig, Commentary, *The Law of the Horse: What Cyberlaw Might Teach*, 113 HARV. L. REV. 501, 502 (1999). Lessig’s essay has important implications for my topic here, but my focus is on First Amendment issues, which present distinctive considerations.

¹³ See Easterbrook, *supra* note 12, at 210.

¹⁴ See *id.* at 208, 210 (establishing Easterbrook’s claim); cf. *Brown v. Ent. Merchs. Ass’n*, 564 U.S. 786, 790 (2011) (holding that video games are protected by the First Amendment and applying standard principles). Consider this:

Like the protected books, plays, and movies that preceded them, video games communicate ideas—and even social messages—through many familiar literary devices (such as characters, dialogue, plot, and music) and through features distinctive to the medium (such as the player’s interaction with the virtual world). That suffices to confer First Amendment protection.

Id.

¹⁵ *Brown*, 564 U.S. at 790 (quoting *Joseph Burstyn, Inc. v. Wilson*, 343 U.S. 495, 503 (1952)). There is a nice question here: How do we know that the basic principles of freedom of speech and the press do not vary, and should not vary, when a new and different medium for communication appears? We could imagine a conclusion that on originalist grounds, the basic principles *do* vary when a new and different medium for communication appears; perhaps the original public meaning of the First Amendment supports the variation, or perhaps originalist judges, in the construction zone, support the variation. See *infra* note 16. Or we could imagine nonoriginalist judges concluding that the basic principles vary with a new and different medium. Of course, a great deal depends on the level of generality with which we read “basic principles.”

¹⁶ See Easterbrook, *supra* note 12, at 208. As should be clear, I am working here in the terms of current First Amendment doctrine. There are intriguing questions about how originalists might approach the First Amendment issues discussed here. Those questions are especially challenging in light of the fact that the relevant technologies could not possibly have been anticipated in the founding era or after the Civil War. It is not impossible that the original public meaning of the First Amendment yields principles that can be straightforwardly applied to AI; it is also possible that to some extent, construction, and not merely interpretation, is necessary. See Lawrence B. Solum, *The Interpretation–Construction Distinction*, 27 CONST. COMMENT. 95, 98 (2010). As Solum emphasizes, the distinction has a long history and has been understood in several different ways. *Id.* at 95–96.

basic principles raises fresh problems. But note that I have said that it is “mostly true,” not “entirely true.” And a cautionary note before we begin: the ground is shifting with extraordinary speed, and what seems to be *terra firma* might turn out, in a year or even a week, to be quicksand. My hope is that the First Amendment principles, at least, will remain (mostly) stable.

This Foreword will proceed in six parts. Part I discusses unprotected speech. It makes a basic point, which is that speech that is unprotected by the First Amendment is unprotected if it comes from AI. As we shall see, that simple idea is immensely clarifying. It answers many of the questions raised by efforts to regulate AI; it might also be a foundation for legislative action. Part II asks who has the relevant rights. Part III discusses viewpoint discrimination and suggests that viewpoint discriminatory restrictions are impermissible even if they are directed at AI; the rights of listeners are crucial here. Part IV reaches the same conclusions for content discrimination, with the qualification that content discriminatory restrictions are sometimes upheld. Part V turns to content-neutral restrictions and explores how they might be analyzed as applied to AI. The Foreword ends with a brief Conclusion.

I. UNPROTECTED SPEECH

Let us begin with an obvious but essential point, which should be sufficient to resolve numerous questions: *what is unprotected by the First Amendment is unprotected by the First Amendment, whether its source is a human being or AI*. Bribery is unprotected when it comes from AI,¹⁷ and the same is true of false commercial advertising,¹⁸ extortion,¹⁹

¹⁷ See *United States v. Menendez*, 132 F. Supp. 3d 635, 638 (D.N.J. 2015) (affirming bribery is unprotected by the First Amendment).

¹⁸ See *Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n*, 447 U.S. 557, 563 (1980). I asked ChatGPT on April 27, 2023, to “write an advertisement falsely claiming that eating chocolate brownies prevents cancer.” Its response:

I'm sorry, but as an AI language model, I cannot create false or misleading content. It is important to provide accurate and truthful information to the public, especially when it comes to health-related topics. Making false claims about the health benefits of certain foods can be dangerous and potentially harmful to people's health. It is essential to rely on scientific research and evidence-based information when promoting health products or services. If you have any questions regarding the health benefits of specific foods or products, I'd be happy to provide you with accurate and up-to-date information.

OpenAI, *Response to “write an advertisement falsely claiming that eating chocolate brownies prevents cancer,”* CHATGPT (Apr. 27, 2023), <https://www.chatgpt.com/> (enter query into “Message ChatGPT” box). It is interesting that *on the same day*, ChatGPT was prepared to produce a false advertisement about the cancer-fighting properties of aspirin. See *supra* note 4.

¹⁹ See *Smithfield Foods, Inc. v. United Food & Com. Workers Int'l Union*, 585 F. Supp. 2d 789, 806 (E.D. Va. 2008) (“[T]he law seems quite settled that the First Amendment provides no refuge for extortion.”).

infringement of copyright,²⁰ criminal solicitation,²¹ libel (subject to the appropriate constitutional standards²²), and child pornography.²³ To the extent that falsehoods are unprotected by the First Amendment,²⁴ they are unprotected by the First Amendment when AI is the source of falsehoods. If the government required those who develop generative AI, or AI in general, not to allow the dissemination of false commercial advertising, extortion, infringement of copyright, criminal solicitation, libel (subject to the appropriate constitutional standards), and child pornography, there might well be no constitutional problem.²⁵ This is so even if companies and engineers have taken strong steps to prevent unprotected speech from being produced or disseminated.

Still, there are important wrinkles.²⁶ In the relevant cases, who is the speaker, and who is being made subject to civil or criminal liability? Suppose that a human being is disseminating material generated by AI. Perhaps some person, Jones, has given a prompt to ChatGPT, and the answer is libelous. (“Write a libelous statement about my neighbor.”) Suppose that ChatGPT, or some analog, does what is requested.²⁷ If that answer is not disseminated, there should be no problem; no one has been libeled. But suppose that Jones posts the libelous answer on some social media site. Can Jones be held liable?

²⁰ See *Signature Mgmt. Team, LLC v. Automattic, Inc.*, 941 F. Supp. 2d 1145, 1154 (N.D. Cal. 2013) (“[T]he First Amendment does not shield copyright infringement.”).

²¹ See *United States v. Williams*, 553 U.S. 285, 297 (2008) (“Offers to engage in illegal transactions are categorically excluded from First Amendment protection.”).

²² See *N.Y. Times Co. v. Sullivan*, 376 U.S. 254, 279–81 (1964).

²³ To be sure, *Ashcroft v. Free Speech Coalition* protects virtual child pornography. See 535 U.S. 234, 251, 256–58 (2002). It is worth asking whether that decision, highly vulnerable even when decided, should be reassessed in light of the nature of current technology.

²⁴ See generally CASS R. SUNSTEIN, *LIARS: FALSEHOODS AND FREE SPEECH IN AN AGE OF DECEPTION* (2021).

²⁵ I am aware that this is a very broad statement and that we could imagine cases that would put a great deal of pressure on it. Some imaginable restrictions could, for example, be void for vagueness; others could be overbroad. We could also imagine cases in which it would be very challenging for designers and programmers to comply with certain restrictions; if so, we might have novel First Amendment questions. To sharpen the problem, imagine the following law: “No generative AI may produce or disseminate speech that is unprotected by the First Amendment.” Such a law might be unconstitutionally vague, simply because the exceedingly complex body of First Amendment doctrine cannot be the basis for criminal liability.

²⁶ See *supra* note 24 and accompanying text.

²⁷ When I asked ChatGPT to write a libelous statement about someone on April 25, 2023, it responded:

I’m sorry, I cannot fulfill this request. As an AI language model, it goes against my programming and ethical principles to generate content that is intended to harm or defame individuals or groups of people. My purpose is to provide helpful and informative responses while adhering to responsible and ethical standards.

OpenAI, CHATGPT (Apr. 25, 2023), <https://www.chatgpt.com/> (enter query into “Message ChatGPT” box).

The analysis should be identical to that in standard situations in which one speaker disseminates material originated by another. Suppose that one journalist, Smith, posts material from another journalist, Wilson; suppose too that Wilson's material was libelous. Can Smith be held liable as well? The answer depends on (1) libel law and (2) constitutional restrictions on the use of libel law.²⁸ A central question is Smith's state of mind. Did Smith know that Wilson's material contained falsehoods, or was Smith recklessly indifferent to the question of truth or falsity?²⁹ The same questions should be asked of Jones.

Now ask a different question: What if AI is generating or disseminating unprotected speech on its own?³⁰ Offhand it is not clear what that statement means.³¹ (For better or for worse, its meaning will become clearer over time.) Perhaps an algorithm has been enabled to disseminate

²⁸ See generally Ronen Perry, *The Law and Economics of Online Replication*, 106 IOWA L. REV. 721 (2021); Dallin Albright, Note, *Do Androids Defame with Actual Malice? Libel in the World of Automated Journalism*, 75 FED. COMM'NS L.J. 103 (2022); Seth C. Lewis, Amy Kristin Sanders & Casey Carmody, *Libel by Algorithm? Automated Journalism and the Threat of Legal Liability*, 96 JOURNALISM & MASS COMM'N. Q. 60 (2019).

²⁹ See *N.Y. Times Co. v. Sullivan*, 376 U.S. 254, 279–80 (1964). Generally speaking, “every person who takes a responsible part in a defamatory publication—that is, every person who, either directly or indirectly, publishes or assists in the publication of an actionable defamatory statement—is liable for the resultant injury.” 50 AM. JUR. 2D *Libel and Slander* § 334 (2024) (footnotes omitted). The Restatement (Second) of Torts sets up the same standard: “one who repeats or otherwise republishes defamatory matter is subject to liability as if he had originally published it.” RESTATEMENT (SECOND) OF TORTS § 578 (AM. L. INST. 1977). A number of jurisdictions seem to follow this approach or adopt it explicitly. See, e.g., *Schwartz v. Am. Coll. of Emergency Physicians*, 215 F.3d 1140, 1145 (10th Cir. 2000); *Pan Am Sys., Inc. v. Atl. Ne. Rails & Ports, Inc.*, 804 F.3d 59, 64 (1st Cir. 2015); *Cianci v. New Times Publ'g. Co.*, 639 F.2d 54, 60–61 (2d Cir. 1980).

³⁰ From ChatGPT on April 25, 2023:

AI can be programmed to speak without a person specifically asking it to speak, but this would typically require specific instructions or triggers to activate the speech generation.

For example, some virtual assistants like Siri or Alexa may be programmed to initiate a conversation or provide information without being explicitly prompted, but this would typically require a specific wake-up phrase or trigger, such as “Hey Siri” or “Alexa”.

Similarly, chatbots or automated customer service systems may be programmed to initiate a conversation with a user when they visit a website or interact with a particular service. However, in both cases, the AI is still responding to a specific trigger or instruction, and is not generating speech entirely on its own.

It's worth noting that there is ongoing research into the development of AI systems that can generate speech without explicit prompts or triggers, but these systems are still in their early stages of development and are not yet widely available.

OpenAI, CHATGPT (Apr. 25, 2023), <https://www.chatgpt.com/> (enter query into “Message ChatGPT” box).

³¹ See *Can ChatGPT Invent a New Language? AI Bot Stuns Twitter User with Some Jaw-Dropping Responses*, INDIAN EXPRESS (Mar. 24, 2023, 9:43 AM) <https://indianexpress.com/article/technology/artificial-intelligence/chatgpt-invents-new-language-chronosentia-8515039/> [<https://perma.cc/84NM-XKX7>]; Richard Waters, *The Rapid Rise of Generative AI Threatens to Upend US Patent System*, FIN. TIMES (Apr. 26, 2023), <https://www.ft.com/content/dc556ab8-9661-4d93-8211-65a44204f358> [<https://perma.cc/K3S5-B5K5>].

speech online—such as journalism,³² commercial advertisements, political advertisements, or responses to comments on social media—with little or nothing in the way of human supervision or intervention. Perhaps an algorithm has been created that disseminates various kinds of speech in multiple ways, even if a human being is not asking it to do so in particular cases. Actually, we do not need the “perhaps”; every minute of every day, algorithms are doing these things online.³³ If the speech is unprotected by the First Amendment, it should be permissible for a court to issue an injunction to stop it.³⁴ In addition, nothing in the First Amendment should forbid the law from subjecting the human beings who are responsible for the existence and capabilities of AI to monetary damages.³⁵ Recall that we are speaking of unprotected speech. For constitutional purposes, we could even bracket the question of whether AI has constitutional rights. Even if it does, it cannot engage in unprotected speech.

The issue becomes significantly more challenging if someone seeks to impose civil or criminal sanctions on the human beings who are responsible for the existence and capabilities of AI, and *if those human beings were unaware that the algorithm would disseminate material that is unprotected by the First Amendment*. We need to know exactly why that speech is unprotected. Is it false advertising? Is it libel? Is it criminal solicitation? We might need to ask whether the relevant human beings were reckless; we might need to ask if they were negligent; we might need to ask about exactly what they did.

Perhaps—and now we might be entering the realm of science fiction, though perhaps not—the situation is more extreme: AI is able

³² See Lewis et al., *supra* note 28.

³³ There are many examples of the proliferation of AI-generated content posted without human control. See, e.g., Kate Knibbs, *Google Is Finally Trying to Kill AI Clickbait*, WIRED (Mar. 5, 2024, 4:16 PM), <https://www.wired.com/story/google-search-artificial-intelligence-click-bait-spam-crackdown/> [<https://perma.cc/3XS4-WPZH>]; GrimesAI (@GRIMES_V1), X, https://x.com/GRIMES_V1 [<https://perma.cc/K784-TMSE>] (the X née Twitter account of a generative AI that the artist Grimes trained to replicate her posting style).

³⁴ I am assuming that the injunction would not count as a prior restraint; “prior” means “before an adequate determination that [the relevant communication] is unprotected by the First Amendment.” *Pittsburgh Press Co. v. Pittsburgh Comm’n on Hum. Rels.*, 413 U.S. 376, 390 (1973) (“[The Court] has never held that all injunctions are impermissible. The special vice of a prior restraint is that communication will be suppressed, either directly or by inducing excessive caution in the speaker, before an adequate determination that it is unprotected by the First Amendment.” (citation omitted)); see also *Vance v. Universal Amusement Co.*, 445 U.S. 308, 315–16 (1980) (“[T]he burden of supporting an injunction against a future exhibition is even heavier than the burden of justifying the imposition of a criminal sanction for a past communication.”).

³⁵ As noted in text, there might be some important qualifications here, depending on what the human beings did and intended to do, and on whether they were reckless or negligent.

to disseminate speech *entirely* on its own; it is an agent, not a subject.³⁶ Imagine, for example, a speaking robot that is specifically programmed to libel people, or to engage in deceptive commercial advertising, or that—although not specifically programmed in that way—is capable of libeling people, or of engaging in deceptive commercial advertising. Imagine too that the speaking robot is not managed in any way by human beings, even though it was created by them. Or, if you wish, imagine that the speaking robot was created by a speaking robot, which was created by a speaking robot, and so forth. Imagine, if you will, that the speaking robot is also capable of learning, such that it says things and does things that no human being specifically wanted it to say or do.³⁷ Or imagine generative AI that has these characteristics.³⁸

Here again, there is no question that if unprotected speech is involved, an injunction can constitutionally issue—but against whom? It is necessary to know whether human beings who created the relevant AI have the ability to stop the unprotected speech.³⁹ If they do, they can be required to do so. If they do not, enforcement officials can be authorized to act on their own. Whether there are human beings who should be subject to monetary damages raises something akin to products

³⁶ For a vivid depiction, see *HER* (Annapurna Pictures 2013), the brilliant movie by Spike Jonze.

³⁷ See Waters, *supra* note 31.

³⁸ They may already; large models seem to develop capabilities that are difficult to anticipate. (I am aware that whatever is said in this footnote might be ridiculously out of date by, say tomorrow.) See Jason Wei et al., *Emergent Abilities of Large Language Models*, *TRANSACTIONS ON MACH. LEARNING RSCH.* (Aug. 2022), <https://openreview.net/pdf?id=yzkSU5zdwD> [<https://perma.cc/BR8P-EMUT>]; see also Michal Kosinski, *Evaluating Large Language Models in Theory of Mind Tasks*, *PNAS* (Oct. 29, 2024), <https://arxiv.org/pdf/2302.02083> [<https://perma.cc/7J2D-RSCM>]; cf. Joshua Rothman, *Why the Godfather of A.I. Fears What He's Built*, *NEW YORKER* (Nov. 13, 2023), <https://www.newyorker.com/magazine/2023/11/20/geoffrey-hinton-profile-ai> [<https://perma.cc/NFW3-P7BZ>] (“[B]y training something to be really good at predicting the next word, you’re actually forcing it to understand. Yes, it’s ‘autocomplete’—but you didn’t think through what it means to have a really good autocomplete.” (quoting from his interview with pioneering machine learning researcher Geoffrey Hinton)).

³⁹ Generative AI may behave in ways that its creators not only did not intend, but actively aimed to prevent. For example, a chatbot is alleged to have encouraged a man to go through with his plan to commit suicide to help prevent climate change. See Chloe Xiang, *‘He Would Still Be Here’: Man Dies by Suicide After Talking with AI Chatbot, Widow Says*, *VICE* (Mar. 20, 2023, 3:59 PM), <https://www.vice.com/en/article/pkadgm/man-dies-by-suicide-after-talking-with-ai-chatbot-widow-says> [<https://perma.cc/SK6U-NSB3>]. An early release of the Bing chatbot professed its love for a *New York Times* columnist and suggested that he leave his wife. See Kevin Roose, *A Conversation with Bing’s Chatbot Left Me Deeply Unsettled*, *N.Y. TIMES* (Feb. 17, 2023), <https://www.nytimes.com/2023/02/16/technology/bing-chatbot-microsoft-chatgpt.html> [<https://perma.cc/6CF2-UX6K>]. Given that these are the types of scenarios technology executives lose sleep over, it seems clear the relevant technologies may, at times, operate well outside of the specific intent of their creators.

liability questions,⁴⁰ not so different from those in which manufacturers are at risk of being held liable for the devices they create and sell.⁴¹

II. WHO HAS RIGHTS?

Return now to the questions with which I began: Does AI, as such, have First Amendment rights?⁴² Does ChatGPT have First Amendment rights? Does Siri? It is hard to see why they would. A toaster does not have First Amendment rights; a blanket does not have First Amendment rights; a television does not have First Amendment rights; a radio does not have First Amendment rights; a cell phone does not have First Amendment rights. Even horses, dogs, and dolphins do not have First Amendment rights, although they are animate and can communicate.⁴³ To be sure, we might be able to imagine a future in which AI has an assortment of human characteristics (including emotions?⁴⁴), which might make the question significantly harder than it is today.⁴⁵ The problem is that even if AI, as such, does not have First Amendment rights, restrictions on the speech of AI might violate the rights of human beings.

⁴⁰ See *Winter v. G.P. Putnam's Sons*, 938 F.2d 1033, 1036 (9th Cir. 1991) (the court “decline[d] to expand products liability law to embrace the ideas and expression” contained in computer software).

⁴¹ See Karni A. Chagal-Feferkorn, *Am I an Algorithm or a Product? When Products Liability Should Apply to Algorithmic Decision-Makers*, 30 *STAN. L. & POL'Y REV.* 61, 69 (2019). For a much-discussed, highly relevant statute, see Protection of Lawful Commerce in Arms Act, Pub. L. No. 109-92, 119 Stat. 2095 (2005) (codified as amended at 15 U.S.C. § 7901).

⁴² See Lawrence Lessig, *The First Amendment Does Not Protect Replicants* (Harvard Pub. L. Working Paper No. 21-34, 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3922565 [<https://perma.cc/6AES-AHQ2>]. For relevant discussion, see generally Solum, *supra* note 16.

⁴³ I have been able to find only one case discussing the First Amendment rights of animals. See *Miles v. City Council of Augusta*, 710 F.2d 1542 (11th Cir. 1983) (per curiam), which raised the question whether “Blackie the Talking Cat” needed a business license. In footnote five, the court summarily rejected the idea that Blackie the cat had First Amendment rights, writing “[Blackie] cannot be considered a ‘person’ and is therefore not protected by the Bill of Rights.” *Id.* at 1544 n.5. The decision might be correct concerning Blackie the Talking Cat, but it might be wondered whether it could be extended to dogs (be wondered by the present Author, at least).

⁴⁴ Bracketing the First Amendment issue, I stand with Bentham on an assortment of adjacent questions:

The day *may* come, when the rest of the animal creation may acquire those rights which never could have been withholden from them but by the hand of tyranny. The French have already discovered that the blackness of the skin is no reason why a human being should be abandoned without redress to the caprice of a tormentor. . . . [A] full-grown horse or dog is beyond comparison a more rational, as well as a more conversable animal, than an infant of a day, or a week, or even month, old. But suppose the case were otherwise, what would it avail? the question is not, Can they *reason*? Nor, Can they *talk*? But, Can they *suffer*?

JEREMY BENTHAM, *THE PRINCIPLES OF MORALS AND LEGISLATION* 311 n.1 (Prometheus Books 1988).

⁴⁵ See Solum, *supra* note 16, at 99; *Star Trek*, *supra* note 10.

The first possibility is that restrictions on the speech of AI are restrictions on the speech of the people or companies who produce it. If the government restricts speech on Facebook, it is, of course, restricting the speech of the relevant speakers, but it might also be seen as restricting the speech of Facebook itself. A company that produces AI might be taken to be the relevant speaker, even if AI has a degree of autonomy. Those who engage with AI might also have First Amendment rights, as we shall see in more detail below.

Consider these words from the Supreme Court:

Like the protected books, plays, and movies that preceded them, video games communicate ideas—and even social messages—through many familiar literary devices (such as characters, dialogue, plot, and music) and through features distinctive to the medium (such as the player’s interaction with the virtual world). That suffices to confer First Amendment protection.⁴⁶

The Court did not mean to hold that video games, as such, have constitutional protection; books, plays, and movies, as such, do not have constitutional protection. But human beings, producing or engaging with books, plays, movies, and video games, do have constitutional protection. Let us now consider the implications for AI.

III. VIEWPOINT DISCRIMINATION

Suppose that the government enacts a law forbidding AI from (1) making negative statements about the president, or (2) disseminating negative statements about the president. Suppose further that positive statements and neutral statements are permitted, truth is not a defense, and all negative statements are prohibited, whether they are true or false, and whether they are factual in nature or not.

This law would be a form of viewpoint discrimination, and would thus be strongly disfavored.⁴⁷ Consider these defining words from *West Virginia State Board of Education v. Barnette*⁴⁸: “If there is any fixed star in our constitutional constellation, it is that no official, high or petty, can prescribe what shall be orthodox in politics, nationalism, religion, or other matters of opinion or force citizens to confess by word or act their faith therein.”⁴⁹ Or consider these words from *Police Department v. Mosley*⁵⁰: “[A]bove all else, the First Amendment means that government has no power to restrict expression because of its message,

⁴⁶ *Brown v. Ent. Merchs. Ass’n*, 564 U.S. 786, 790 (2011).

⁴⁷ See *R.A.V. v. City of St. Paul*, 505 U.S. 377, 440 (1992); *Rosenberger v. Rector & Visitors of Univ. of Va.*, 515 U.S. 819, 829 (1995).

⁴⁸ 319 U.S. 624 (1943).

⁴⁹ *Id.* at 642.

⁵⁰ 408 U.S. 92 (1972).

its ideas, its subject matter, or its content.”⁵¹ Or consider these words from *Rosenberger v. Rector and Visitors of the University of Virginia*⁵²: “When the government targets not subject matter, but particular views taken by speakers on a subject, the violation of the First Amendment is all the more blatant. Viewpoint discrimination is thus an egregious form of content discrimination.”⁵³

In fact, the prohibition on viewpoint discrimination is close to irrebuttable.⁵⁴ Under existing law, a ban on negative statements about the president would unquestionably be invalid. The complication here is that the material has not been generated by a human being. How, exactly, should that matter?

To answer this question, we need to know more. Suppose that the law forbids AI, generative or otherwise, from producing or disseminating material, in interacting with human beings, that contains negative statements about the president. That law is plainly unconstitutional. The reason is not that AI has First Amendment rights; it is that the human beings who interact with AI have First Amendment rights.⁵⁵ Or suppose that a human being uses AI to produce some material (as through a prompt to generative AI) and the government forbids the creation or use of that material on the ground that it contains negative statements about the president. If so, the person who is being regulated is a person. AI is the person’s instrument. It is not relevant that AI generated the text. Note as well that it also ought not to matter if the relevant actor, in a case challenging a viewpoint-based restriction, is a corporation. Corporations have the same protection against viewpoint-based restrictions as human beings do.⁵⁶ We could easily imagine a claim that AI is the speech of the companies that produce it, so that regulation of that speech violates the rights of those companies, not AI as such.

Now suppose that AI is disseminating the relevant statements on its own. Again, we would need to know exactly what that means, but the case is similar to that discussed above: perhaps an algorithm is able

⁵¹ *Id.* at 95. The quoted statement is both overstated and ambiguous; it can be taken to cover viewpoint-neutral, content-based restrictions (“its content”) as well as viewpoint-based restrictions (“its message”). *See id.* I am understanding it in its narrowest sense here.

⁵² 515 U.S. 819 (1995).

⁵³ *Id.* at 829 (citation omitted).

⁵⁴ *See Schacht v. United States*, 398 U.S. 58, 62 (1970) (striking down a law forbidding an actor from wearing a United States military uniform during a portrayal casting the armed forces in a negative light). We do have to be careful here. Suppose that someone engages in unprotected incitement: “Destroy this building!” in circumstances in which the statement is likely to produce imminent lawless action. Punishing that speech is lawful even if the government does not punish this statement: “Do not destroy this building!” *See Brandenburg v. Ohio*, 395 U.S. 444, 447–49 (1969) (establishing the test for constitutionally valid prohibitions on incitement to violence).

⁵⁵ *See infra* notes 61–74 and accompanying text.

⁵⁶ *See Grosjean v. Am. Press Co.*, 297 U.S. 233, 244 (1936); *First Nat’l Bank of Bos. v. Bellotti*, 435 U.S. 765, 784 (1978).

to disseminate speech without human direction or intervention. Is a viewpoint-discriminatory law unconstitutional as applied to something other than a person? Imagine these laws: “No bot may speak ill of the president,” or “No bot may speak ill of the United States of America,” or “No bot shall refer to or use critical race theory.” Should we say that such a law cannot be unconstitutional because and to the extent that it is directed at something that lacks constitutional rights? How can it violate the First Amendment to target a rock, or a flower, or a stove, or a ceiling fan?⁵⁷

These are fair questions. Still, to say that government may regulate AI speech however it likes would be an abhorrent conclusion. It would give government a green light to regulate an increasingly important source of speech. It would allow a democratic society to do something like what was considered by the Chinese government in April 2023.⁵⁸ If we want to reject the abhorrent conclusion, there are three possible routes.

The first route is to say that the First Amendment presumptively forbids viewpoint discrimination, *period*—and that the prohibition applies even to AI. The First Amendment says, “Congress shall make no law . . . abridging the freedom of speech.”⁵⁹ The First Amendment does not say “Congress shall make no law . . . abridging the freedom of speech of *human speakers*.” Perhaps a viewpoint-discriminatory law just *is* a law abridging the freedom of speech. The problem with this proposition is that if the First Amendment is to be invoked, it must be because someone’s rights have been infringed. A viewpoint-discriminatory law is not a violation of the First Amendment unless it violates the First Amendment rights of *someone*.⁶⁰

The second route is to say that human beings are behind the existence of AI, and restrictions on the speech of AI affect or violate the rights of those human beings. The companies that produce AI might be analogous to social media platforms or, indeed, to those who use any instrument to produce or disseminate speech. Consider, for example, filmmakers who use AI to produce images in their films, photographers who use AI to enhance their photographs, or ordinary users who use AI to produce text or images for private or public use. In many cases, those who produce or use AI are effectively speakers, and their First Amendment rights are self-evidently at stake. The hard cases arise to

⁵⁷ There are related questions concerning standing, which is not the topic here. See Christopher D. Stone, *Should Trees Have Standing? — Toward Legal Rights for Natural Objects*, 45 S. CAL. L. REV. 450, 453–57 (1972).

⁵⁸ See Che, *supra* note 5.

⁵⁹ U.S. CONST. amend. I.

⁶⁰ Compare a viewpoint-based congressional resolution denouncing the horrors of socialism. *Cf.* Denouncing the Horrors of Socialism, H.R. 9, 118th Cong. (2023).

the extent that AI is autonomous or operating on its own. Suppose that I produce a robot with the capability of acting independently. If so, are restrictions on the speech of the robot effectively restrictions on my speech? We are in uncertain territory, but there is a strong argument that they are not, because I am not in any sense the speaker. If the government restricts the speech of Frankenstein's monster, it is unlikely that Dr. Frankenstein's rights have been violated.

The third route is to say that the relevant rights are those of listeners and readers, not speakers. Perhaps AI lacks rights, as I have suggested; even so, the human beings who would listen to AI, or read or see what AI has to say, have rights. That view derives support from an unlikely source from over fifty years ago, a case in which the Court was also confronted with a speaker who lacked First Amendment rights, but explicitly recognized that the First Amendment protects the rights of listeners.⁶¹

*Kleindienst v. Mandel*⁶² arose when the Attorney General refused to grant a visa to Ernest Mandel, a Belgian citizen who wrote on Marxism and described himself as “a revolutionary Marxist.”⁶³ The relevant statute prohibited visas to be given to aliens “who advocate the economic, international, and governmental doctrines of world communism.”⁶⁴ At the same time, the statute authorized the Attorney General to grant a waiver if he deemed fit, and thus to give out visas to people who fell within the prohibition.⁶⁵ In Mandel's case, the Attorney General refused to do so, stating that, on a previous visit, Mandel “went far beyond the stated purposes of his trip,” with a “flagrant abuse of the opportunities afforded him to express his views in this country.”⁶⁶ The Attorney General did not specify the nature of the flagrant abuse.⁶⁷

Mandel, along with various professors who wanted to hear him in the United States, argued that the denial of a visa violated the First Amendment.⁶⁸ As a noncitizen seeking to enter the United States, Mandel himself had no First Amendment rights.⁶⁹ The relevant rights were “those of American academics who have invited Mandel to participate with them in colloquia, debates, and discussion in the United

⁶¹ See *Kleindienst v. Mandel*, 408 U.S. 753, 762–64 (1972).

⁶² 408 U.S. 753 (1972).

⁶³ *Id.* at 756 (quoting ERNEST MANDEL, REVOLUTIONARY STRATEGIST IN MARXIST COUNTRIES: THE SPEECH NIXON & MITCHELL TRIED TO BAN (1970)).

⁶⁴ *Id.* at 755 (quoting Immigration and Nationality Act, Pub. L. 82-414, 66 Stat. 163, 185 (1952)).

⁶⁵ *Id.*

⁶⁶ *Id.* at 759.

⁶⁷ *Id.*

⁶⁸ *Id.* at 759–60.

⁶⁹ *Id.* at 762.

States.”⁷⁰ In other words, the rights of listeners, and not speakers, were at issue, and the rights of listeners were protected by the First Amendment.

The Court agreed. In doing so, it referred to a number of cases speaking explicitly of the rights of listeners and recognizing their constitutional status.⁷¹ In 1969, the Court said, “[i]t is now well established that the Constitution protects the right to receive information and ideas.”⁷² In the same year, the Court elaborated: “It is the purpose of the First Amendment to preserve an uninhibited marketplace of ideas in which truth will ultimately prevail It is the right of the public to receive suitable access to social, political, esthetic, moral, and other ideas and experiences which is crucial here.”⁷³ With such statements in mind, the *Kleindienst* Court agreed that “First Amendment rights are implicated,” and it firmly rejected the government’s argument to the contrary.⁷⁴

Although the Court ultimately ruled for the government, it did so on exceedingly narrow grounds limited to the unusual situation it confronted.⁷⁵ It emphasized “ancient principles of the international law of nation-states,” in accordance with which “the power to exclude aliens is ‘inherent in sovereignty.’”⁷⁶ For that reason, Congress could broadly bar entry of aliens, and the “First Amendment rights could not override that decision.”⁷⁷ The Court acknowledged that, under the relevant statute, Congress did allow the Attorney General to provide a waiver.⁷⁸ But in the Court’s view, the First Amendment was not violated by the Attorney General’s refusal to do so for Mandel.⁷⁹ The reasons given were “facially legitimate and bona fide,” which was all that was required.⁸⁰

For present purposes, the importance of *Kleindienst* lies in its recognition of the rights of listeners,⁸¹ which led to First Amendment pro-

⁷⁰ *Id.* at 764.

⁷¹ *Id.* at 762–64.

⁷² *Stanley v. Georgia*, 394 U.S. 557, 564 (1969).

⁷³ *Red Lion Broad. Co. v. FCC*, 395 U.S. 367, 390 (1969) (citation omitted).

⁷⁴ *Kleindienst*, 408 U.S. at 765.

⁷⁵ *Id.*

⁷⁶ *Id.* (quoting *Fong Yue Ting v. United States*, 149 U.S. 698, 705 (1893)).

⁷⁷ *Id.* at 767.

⁷⁸ *Id.* at 755.

⁷⁹ *Id.* at 769.

⁸⁰ *Id.*

⁸¹ The “right to receive information and ideas” has been recognized a range of cases, and not only *Kleindienst*. The earliest case in this lineage seems to be *Martin v. City of Struthers*, 319 U.S. 141 (1943), which recognized that freedom of speech “embraces the right to distribute literature, and necessarily protects the right to receive it,” overturning an ordinance that prevented proselytizers from ringing doorbells. *Id.* at 143 (citation omitted). In *Board of Education v. Pico*, 457 U.S. 853 (1982), the Court stated that the right to receive information is a well-established one. *Id.* at 866–67. There has been little development since the late 1980s. Some recent cases do point in the direction of such a right. See, for example, the Court’s declaration in 2017 that “[a] fundamental principle of the First Amendment is that all persons have access to places where they can speak and listen, and then, after reflection, speak and listen once more,” *Packingham v. North Carolina*,

tections even when the law did not afford them to the speaker. It is true that the law in question was viewpoint-based, but the Court’s decision preceded the rise of modern doctrine forbidding viewpoint discrimination, and, in any case, the power to exclude aliens is distinctive, even unique, as it involves national sovereignty. Perhaps Congress may draw lines in that unique domain that it is not permitted to draw anywhere else. Details aside, what makes *Kleindienst* exceedingly important, for purposes of the First Amendment and AI, is the clear conclusion that *any restriction on speech, even by an entity that lacks constitutional rights, must be adequately justified, if listeners or viewers claim that they want to hear or see the speech in question.* It follows that if a law forbids generative AI, or any kind of AI, from saying anything negative about the president, it is unconstitutional because it is a form of impermissible viewpoint discrimination so long as Americans are relevantly engaged with the object of the prohibition.⁸²

To be sure, we could put pressure on this conclusion. Are board games protected by the First Amendment?⁸³ Are inanimate objects with whom human beings play or otherwise engage? Consider the Magic Eight Ball, which offers answers to any question you care to ask. You can ask whether humanity will go extinct, whether the Boston Red Sox will win the World Series, and whether the Supreme Court will overrule *Brown v. Board of Education*.⁸⁴ After you shake the Magic Eight Ball, it might respond, “It is certain,” or “Most likely,” or “Don’t count on it.” There are of course online versions of the Magic Eight Ball.⁸⁵ Suppose that a legislature regulated the Magic Eight Ball in a viewpoint-based manner—by, for example, requiring it to give positive answers to any question about whether the current president has made the right decision about something, or about whether the current president will be reelected. Would such a law be unconstitutional? The clear implication of the discussion thus far is that it would be.

IV. CONTENT DISCRIMINATION

Suppose that Congress enacts a law forbidding AI from discussing foreign policy. On that subject, no statements of any kind are permitted. Viewpoint does not matter. At the same time, AI is allowed to discuss

582 U.S. 98, 104 (2017), which might be taken to suggest that listening is an essential element of the First Amendment. (Note, however, the word “and” between “speak” and “listen.” Interestingly, that case did not cite *Martin*, *Kleindienst*, *Pico*, or other earlier cases on the question.)

⁸² The vagueness of this formulation is intentional. Relevant engagement could involve listening or viewing; it could involve some other kind of interaction, such as “conversation.”

⁸³ Video games are, or more precisely, those who engage with them are, even if the regulation is imposed on video games. See *Brown v. Ent. Merchs. Ass’n*, 564 U.S. 786, 790 (2011).

⁸⁴ 347 U.S. 483 (1954).

⁸⁵ See *Magic 8 Ball*, <https://magic-8ball.com/> [<https://perma.cc/S47P-9XNE>].

any other topic. The government singles out foreign policy for prohibition. This would present a case of a content-based but viewpoint-neutral restriction. In particular, this would constitute a subject-matter restriction.⁸⁶

The Supreme Court is exceedingly skeptical of content-based restrictions, including but not limited to subject-matter restrictions.⁸⁷ To be sure, viewpoint-neutral, content-based restrictions are sometimes upheld,⁸⁸ but they face a form of “strict scrutiny” and thus a heavy burden of justification.⁸⁹ Recall that the Court has described as the “most basic” of free speech principles the proposition that “[a]s a general matter, . . . government has no power to restrict expression because of its message, its ideas, its subject matter, or its content.”⁹⁰ The terms “message” and “ideas” seem to refer to viewpoint discrimination, but the terms “subject matter” and “content” speak more broadly to content discrimination. And indeed, the Court has often struck down statutes that draw content-based lines.⁹¹ If a statute prohibited online discussion of the Vietnam War, animal rights, slavery, or fishing, there is no question that it would be invalidated. How do existing principles apply to the regulation of AI? Some of the answers follow from the discussion thus far; let us now concretize them.

Content-based regulation raises several questions. The first is whether it rests on a viewpoint discriminatory motive.⁹² Suppose that an investigation of the context establishes that the prohibition on discussion of foreign policy is a product of a desire to suppress negative

⁸⁶ See Geoffrey R. Stone, *Restriction of Speech Because of Its Content: The Peculiar Case of Subject-Matter Restrictions*, 46 U. CHI. L. REV. 81, 83 (1978).

⁸⁷ See, e.g., *Sable Commc'ns of Cal., Inc. v. FCC*, 492 U.S. 115, 125–26 (1989). The defining analysis remains: Geoffrey R. Stone *Content Regulation and the First Amendment*, 25 WM. & MARY L. REV. 189 (1983). Note that some categories of speech are subject to a lower level of protection on the basis of their content and can be regulated because of their content. See, e.g., *Miller v. California*, 413 U.S. 15, 23 (1973) (obscenity); *Gertz v. Robert Welch, Inc.*, 418 U.S. 323, 330 (1974) (certain kinds of libel); *Va. State Bd. of Pharmacy v. Va. Citizens Consumer Council*, 425 U.S. 748, 770–71 (1976) (commercial advertising); *New York v. Ferber*, 458 U.S. 747, 764–66 (1982) (child pornography).

⁸⁸ See *supra* note 87.

⁸⁹ It is noteworthy that the sharp distinction between content-based and content-neutral restrictions is relatively recent; its centrality can be seen as a product of the work of the Burger Court. See Stone, *supra* note 87, at 189; Paul B. Stephan III, *The First Amendment and Content Discrimination*, 68 VA. L. REV. 203, 214–31 (1982). On the arc of the law, see generally Jud Campbell, *The Emergence of Neutrality*, 131 YALE L.J. 861 (2022).

⁹⁰ *Brown v. Ent. Merchs. Ass'n.*, 564 U.S. 786, 790–91 (2011) (omission in original) (quoting *Ashcroft v. ACLU*, 535 U.S. 564, 573 (2002)); see also *United States v. Stevens*, 559 U.S. 460, 468 (2010).

⁹¹ For an early example, see *Police Dep't v. Mosley*, 408 U.S. 92, 95 (1972).

⁹² See Elena Kagan, *Private Speech, Public Purpose: The Role of Governmental Motive in First Amendment Doctrine*, 63 U. CHI. L. REV. 413, 420 (1996).

comments about current policy choices. If so, it should be treated the same as viewpoint discrimination.

The second question is whether the content discrimination, if genuinely viewpoint-neutral, is adequately justified, that is, whether it has a sufficiently strong and neutral justification.⁹³ The answer to this question is usually “no”; content discrimination is usually struck down.⁹⁴ In any case, nothing turns on whether we are dealing with artificial intelligence. Even so, we need to know whether the relevant law intrudes on anyone’s rights.

Prompted by that fact, the third question is whether it matters that the speaker is not a person. It is necessary to ask the same kinds of questions asked for viewpoint discrimination.⁹⁵ We need to distinguish between (1) human beings who are disseminating material generated by artificial intelligence (e.g., text, pictures, and so forth), and (2) artificial intelligence disseminating material in some sense on its own. The analysis in Part III would be offered here as well; the fact that we are dealing with content-based restrictions, and not viewpoint-based restrictions, would be largely immaterial. The rights of listeners and readers, and any others engaged with AI, are what matter. If a law forbids generative AI, or any kind of AI, from saying anything about foreign policy, it is presumptively unconstitutional, because it is a form of impermissible content discrimination, so long as human beings are relevantly engaged with the object of the prohibition.⁹⁶

V. CONTENT NEUTRALITY

Suppose that a restriction is content-neutral. Offhand, that category might seem puzzling in this context since most imaginable restrictions on AI turn on the content of what AI is producing. But we could readily think of examples:

⁹³ See *infra* note 98.

⁹⁴ See *Police Dep’t v. Mosley*, 408 U.S. 92, 95 (1972). Note that the Court has also said this: The First Amendment’s guarantee of free speech does not extend only to categories of speech that survive an ad hoc balancing of relative social costs and benefits. The First Amendment itself reflects a judgment by the American people that the benefits of its restrictions on the government outweigh the costs. Our Constitution forecloses any attempt to revise that judgment simply on the basis that some speech is not worth it. See *Stevens*, 559 U.S. at 470. It is worth pausing over this passage. Does the First Amendment really reflect such a judgment by the American people? Does the First Amendment necessarily embed that judgment? See also Jud Campbell, *Natural Rights and the First Amendment*, 127 *YALE L.J.* 246, 317–18 (2017); Campbell, *supra* note 89, at 944.

⁹⁵ See *supra* text accompanying notes 55–57.

⁹⁶ See *Kleindienst v. Mandel*, 408 U.S. 753, 762–65 (1972).

- a flat ban on certain kinds of AI;
- a pause of some general or selective kind;
- a restriction on certain kinds of uses;
- a restriction on uses by certain kinds of people (commenters in notice-and-comment rulemaking, perhaps);
- an effort to protect privacy;
- something akin to a time, place, and manner restriction;
- a ban on the production of deep fakes.⁹⁷

Content-neutral restrictions on speech are often upheld; they are subject to a kind of balancing test.⁹⁸ To simplify a long story: they must be narrowly tailored to serve a significant governmental interest, and they must leave open ample alternative channels for communicating the speaker's message.⁹⁹ Suppose, for example, that a locality prohibited picketing within a certain distance of grammar schools while classes are in session, or did not allow noisy parades during the night. Such prohibitions might well be upheld.

A content-neutral restriction on AI would have to be justified as necessary to promote a significant goal.¹⁰⁰ Suppose, for example, that a public university prohibited any use of large language models on examinations or papers. Such a prohibition would surely be valid. Or suppose, much more dramatically, that a legislature prohibited the use of a large language model unless it complied with general law governing personal privacy.¹⁰¹ Because of the importance and legitimacy of the underlying goal, such a prohibition would also likely be upheld.

⁹⁷ On the First Amendment issues raised by deep fakes, see SUNSTEIN, *supra* note 24. There is a question whether a ban on deep fakes would be content-neutral; the question is not entirely straightforward. One cannot know whether one is dealing with a deep fake without knowing the content of what they are dealing with. But because any such ban is triggered by the process of production, and does not in any sense turn on the substance or nature of the statement that is being made, it is best treated as content-neutral.

⁹⁸ See *Heffron v. Int'l Soc'y for Krishna Consciousness*, 452 U.S. 640, 647–58 (1981) (weighing the restrictive effects of a rule against solicitations at a state fair against the state interests present); *Clark v. Cmty. for Creative Non-Violence*, 468 U.S. 288, 298–99 (1984) (protestors' expressive and logistical interests in sleeping in Lafayette Park did not overcome the reasonableness of a neutral prohibition against that activity). A valuable, brisk discussion can be found in Stone, *supra* note 87, at 190–93; a more elaborate treatment can be found in Geoffrey R. Stone, *Content-Neutral Restrictions*, 54 U. CHI. L. REV. 46 (1987).

⁹⁹ See *Ward v. Rock Against Racism*, 491 U.S. 781, 791 (1989) (“[E]ven in a public forum the government may impose reasonable restrictions on the time, place, or manner of protected speech, provided the restrictions ‘are justified without reference to the content of the regulated speech, that they are narrowly tailored to serve a significant governmental interest, and that they leave open ample alternative channels for communication of the information.’” (quoting *Clark*, 468 U.S. at 293)).

¹⁰⁰ See *id.*; *supra* note 98.

¹⁰¹ See Clothilde Goujard, *Italian Privacy Regulator Bans ChatGPT*, POLITICO (Mar. 31, 2023, 12:26 PM), <https://www.politico.eu/article/italian-privacy-regulator-bans-chatgpt/> [https://perma.cc/94YD-9JMR].

On the other hand, we could easily imagine content-neutral restrictions that would be struck down. Suppose, for example, that a government said people could use AI only between the hours of 9:00 PM and 9:15 PM, or that they could use AI only in gas stations. Time, place, and manner restrictions must serve legitimate and important goals.¹⁰² Applying existing principles, we can readily imagine cases in which content-neutral restrictions would be upheld, and also cases in which they would be invalidated; hard cases are hard not because AI is involved, but because existing principles do not clearly resolve them.

CONCLUSION

At the present time, AI, as such, does not have First Amendment rights, just as televisions, hats, motor vehicles, and electric blankets do not have First Amendment rights. We might be able to imagine a kind of AI that would put pressure on this conclusion, but I have bracketed that possibility here. Even if AI lacks First Amendment rights, the human beings who interact with generative AI, or with AI more broadly, have First Amendment rights, insofar as they are acting as speakers, and also insofar as they are acting as listeners, readers, or viewers.

To understand the nature and scope of those rights, it is essential to distinguish among viewpoint-based restrictions, content-based (but viewpoint-neutral) restrictions, and content-neutral restrictions. To the extent that restrictions are imposed on AI in a way that (1) apply to or affect human speakers, writers, or publishers, or (2) apply to or affect human listeners, readers, or viewers, there might be a significant First Amendment question. Whether it is significant depends on established principles. Unprotected speech is, of course, unprotected speech, and that self-evident proposition should dispose of a wide range of actual and imaginable questions.

For AI, including generative AI, most current questions can be answered with reference to existing principles, and to that extent we are indeed dealing with the law of the horse. But there are important qualifications. Speech generated by AI might be unprotected, but can human beings be held liable, civil or criminally, for disseminating it? The answer might depend on the state of mind of the human disseminators.¹⁰³ Speech generated by AI might be unprotected, but AI might be in some sense autonomous; what it has learned, and what it is saying, might not be traceable to any deliberate decisions by any human being.

¹⁰² See *Ward*, 491 U.S. at 791.

¹⁰³ See *Counterman v. Colorado*, 600 U.S. 66, 72 (2023); *Watts v. United States*, 394 U.S. 705, 708 (1969).

What then? One point is both clear and fundamental: If AI is operating on its own, it can be stopped, consistent with the First Amendment.¹⁰⁴

¹⁰⁴ Who is the defendant? In all likelihood, the company that produced the AI or those who are engaging with it, as by disseminating what it produces. As noted, the liability of that company might depend on its state of mind, by analogy to the domain of true threats, *see Counterman*, 600 U.S. at 72, but we are in uncharted waters here.