

Artificial Intelligence Regulation Threatens Free Expression

BY DAVID INSERRA

The past several years have seen rapid technological advancements in artificial intelligence (AI). As exemplified by ChatGPT, new generative AI technologies have gained widespread public usage in various domains, including productivity, art, education, commerce, and science. By performing complex tasks previously exclusive to humans, such as developing artistic imagery or synthesizing arguments, these technologies have the potential to radically change how we interact with the world as citizens, workers, consumers, creators, advocates, families, and individuals. Generative AI specifically has significant and growing implications for expression since it enables users to better express themselves and gain a deeper understanding of the world and the perspectives of others.

However, due to the massive potential of this technology, concerns about AI have also grown.¹ These concerns are often driven by worries of employment displacement caused by automation or existential dread inspired by science fiction scenarios such as the machines of *The Matrix*

or Skynet in *The Terminator*, where the machines take over. Nevertheless, similar to concerns about previous technologies, fears over AI have reached the point of a moral panic. As with other technologies that have expanded speech, this is especially true regarding the fear that AI technology will advance “harmful” ideas and beliefs. To stop the spread of hate speech and misinformation, governments and technology companies are increasingly attempting to impose restrictions—many of which are drawn from the debate over social media content moderation—that would restrain expression and innovation. Unfortunately, much of what is being proposed is hostile to various forms of expression that are mainstream and lawful.²

With AI set to impact so much of our lives, the decisions over what kinds of speech AI should or should not be allowed to generate will likely become far greater battles than those fought over social media content policies. For many users, the disastrous release of Google’s Gemini AI, which manifested clear ideological requirements in its coding, raised these concerns. While a private company has



DAVID INSERRA is a fellow for free expression and technology at the Cato Institute.

every right to develop products that present its viewpoints and biases, users are free to leave products that don't meet their needs and choose alternative products that serve them more effectively.

The most significant threats to the expressive power of AI are government mandates and restrictions on innovation. These threats can take the form of broad regulations that generally limit innovation and prevent new companies from entering the market, or they can take the form of efforts to curtail specific types of AI products or expression in the name of safety, responsibility, or alignment with certain ethics or values. Beyond the economic and national security reasons to empower AI development, AI also has the potential to be a tool to expand free expression, as many previous technologies have. But its potential as a tool for greater free expression requires policymakers to embrace innovation and allow citizens to use this new technology freely.

WHAT IS AI?

AI has been defined in a variety of ways. An October 2023 executive order in the United States described it as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.”³ Other definitions of AI focus on the “human” nature of the work that it can perform or its simulation of “human intelligence and problem-solving capabilities.”⁴ While AI may seem like a new phenomenon to many, older versions of AI have been with us for as long as we have had devices with significant computing capabilities. Serious research into AI began in the 1950s and led to the creation of tools such as spam filters and gaming machines such as the chess computer Deep Blue, which defeated world chess champion Garry Kasparov in 1997. But it is also true that there is something different about new forms of AI—these new AI models have dramatically advanced in their performance due to greater computational power that can handle larger datasets and innovations in how effectively machines can learn from those datasets.⁵

Indeed, the generative AI models that have emerged, such as ChatGPT, along with other related AI tools, such as conversational AI and predictive AI, possess vastly superior abilities to generate written products, images, and conversations and to produce more advanced analyses

and predictions.⁶ Current AI tools have a wide number of applications that society is only beginning to discover. Such AI tools can support:

- artistic and entertainment fields through audio and image generation;
- social media content moderation at scale;
- educational efforts through tailored assignments and learning aids;
- customer service through effective virtual assistants and chatbots;
- communication across language barriers through translation;
- scientific research through parsing large amounts of data, such as genetic sequences or complex physics computations;
- health care through assisting with medical administrative tasks and providing diagnoses;
- any number of computer programs that write computer code; and
- the operation of heavy machinery.

While these applications are expansive, the unique opportunities for expression are among the most notable. As a February 2024 report from the Future of Free Speech, an organization dedicated to advancing a culture of free expression around the world, says on AI:

Gen AI can empower freedom of expression as never before, supercharging the already exponential growth to impart and access information and ideas launched by the internet. For the first time in history, convincing human-sounding content can be generated algorithmically.⁷

AI can provide low-cost, high-quality communication tools to individuals and organizations looking to speak their minds and share their views. Politicians and activists can effectively reach new audiences through translation tools or by using AI to customize speech. About one-third of all generative AI users in September 2023 were using AI tools to learn about various topics.⁸ And the use of AI will only grow as the technology expands and users increasingly understand and adopt it.

NEW TECHNOLOGY, SAME MORAL PANIC

As with all new disruptive technologies, especially with new expressive technology, AI challenges the status quo and creates various worries and concerns. The world currently finds itself amid what many have called a moral or elite panic over the potential abuses or doomsdays that AI could herald.⁹ Just the term “AI” brings to mind the existential threat posed by sentient machines that destroy humanity conveyed in popular art and culture.¹⁰ Many hold a less catastrophic worry similar to that of the Luddites in early 19th-century England or carriage makers at the start of the 20th century when automobiles were being produced. They fear that AI will take away jobs and leave multitudes of workers unemployed. Others worry about the ways criminals and terrorists might abuse AI to further their illegal activities. But bad actors have always used any new technology to advance their own ends.

Perhaps the greatest panic over AI is about its expressive and informational elements and how these can be used to advance speech and viewpoints designated as harmful. The World Economic Forum’s *Global Risks Report* listed AI-powered misinformation and disinformation as the most severe threat to the world in the next two years.¹¹ Even in the long term, this report rates AI-powered misinformation and “adverse outcomes of AI technologies” as the fifth- and sixth-greatest risks, respectively. The report reached this conclusion based on the input of “1,490 experts across academia, business, government, the international community and civil society.” Elites and experts are truly worried about the power of AI and online speech.

The leading researchers and tech companies in the AI field also fear AI spreading harmful content. The attention given to Gemini’s high-profile failures is an obvious example. Gemini’s near-universal refusal to generate images of white people, even when specifically asked to generate images of white people or when asked to generate historically accurate images of white individuals, is the result of choices made by Google to favor ideology over accuracy. The paper that Google put out describing how Gemini operates and is trained shows multiple levels of ideological preferences for various diversity considerations as well as the imperative to prevent “harm-inducing” hate speech and other biased results.¹² Naturally, a company might want to implement

safeguards to prevent its products from being grossly abused, as seen in several cases where AI tools adopted prejudicial viewpoints.¹³ But reports on Gemini state that Google knew that its product was overindexing to prevent outcomes its instructions deemed harmful and nondiverse.¹⁴

If we look beyond Google on this issue, we can see that many researchers follow a similar trend in indexing against content they define as harmful.¹⁵ Major AI companies make countless references to their AIs’ safety, responsibility, and alignment with human values and ethics.¹⁶ Certainly, firms should prevent their AIs from generating content that poses true safety and ethical risks, such as malware or instructions for bioweapons, but beyond those more apparent harms, what does it mean to have an AI “aligned with human intentions and values”?¹⁷ While such statements are vaguely laudable, they beg some serious questions: Which values are human values, and which are not? Who is deciding what is and is not a human value? While these companies have the right to do what they want with their products, they are seizing a fairly broad mandate to set moral and ethical norms—those informal values and standards that society uses to self-regulate and guide acceptable conduct.

Unfortunately, the norms of free expression being adopted in AI are not the same as those that govern significant portions of liberal democracies today. The Future of Free Speech recently analyzed the policies of the largest public AIs and found that “Gen AI providers seem to have opted for a sanitized model straight away, ignoring or minimizing freedom of expression considerations . . . even though—unlike social media posts—the output of chatbots is not automatically disseminated to the public.”¹⁸

As alluded to in the report, this battle over the rules and norms of AI-enabled expression springs from the ongoing and hotly debated issue of social media content moderation. Many social media companies and the broader trust and safety field—the researchers, activists, and practitioners of online content moderation—have advocated for or implemented a narrower view of acceptable online speech than constitutional protections or liberal norms would affirm. Some of this is due to simple business realities—users and advertisers don’t want social media feeds filled with spam and racist speech. But increasingly often, social media companies are adopting and helping to set norms that view growing amounts of controversial social and political speech as not just offensive

or incorrect but actually harmful.¹⁹ While these firms have the right to set their own standards on what is harmful and not allowed on their platforms, they derive these standards from academia, activists, and government actors generally hostile to free expression.²⁰

Other governments frequently regulate speech to the detriment of free expression since no other country has the same legal protections for expression from government intervention as the United States. The US Supreme Court recently decided against social media users and states that had sued the government for exerting informal pressure on platforms into moderating the users' content. The Court found that "while the record reflects that the government defendants played a role in at least some of the platforms' moderation choices," the user and states were not able to clearly establish standing.²¹ More concretely, the US government, through various agencies, has attempted to suppress what it believes to be harmful speech. The Department of Homeland Security attempted to establish its Disinformation Governance Board, which is now widely criticized.²² The State Department's Global Engagement Center directly supported and funded efforts to prevent American media organizations from advertising online.²³ The National Science Foundation is funding the efforts of academics and researchers to develop AI tools that social media companies can use to combat various forms of disfavored speech, such as misinformation and hate speech.²⁴ Significant academic and advocacy work has been conducted to research and stop various forms of harmful speech, including massive fact-checking enterprises and interest groups pressuring social media companies to remove speech disliked by certain groups.²⁵

But as contentious as these battles are, the applications of AI are so vast and transformative that the rules governing it, whether they are policies set by private companies or government regulations, may have a far greater impact on the future of free expression and inquiry than the outcomes of these battles. It is therefore deeply concerning to see tech companies—which have done so much to practically expand users' speech and access to information—borrowing from social media norms and developing even more speech-restrictive AI policies, as such restrictions limit important social and political conversations and viewpoints. By integrating speech-restrictive norms widely across new

disruptive technologies, our societies are effectively erecting new speech codes that not only apply to technological services and products but also create broad cultural expectations and beliefs critical of free expression, both online and offline.

The market for AI tools, however, has an important distinction from social media that can encourage greater expression—the absence of network effects. Access to a large network on a social media platform is often a benefit—for example, influential social media users can reach many other users with their message. But if users dislike aspects of a given platform, it can be difficult, though certainly not impossible, to leave that network. AI tools, however, can be changed as quickly as one can sign up for an account with an alternative provider. As a result, the speech-restrictive norms and policies currently being established among large AI firms have less staying power. Anyone using an AI product who is angered or disappointed by the experience is mere clicks away from an alternative product. As current AI products become more biased or restrictive of speech, the market for alternative products that function on different values and perspectives becomes more abundant. Increasingly individualized and personalized AI tools are being developed and sold, making it likely that soon all users could have access to personalized AIs that fit their unique values and needs.²⁶

GOVERNMENT REGULATION IS THE GREATEST THREAT TO AI-ENABLED EXPRESSION

Given that increasingly capable, diverse, and personalized AIs are a likely market response to consumer demands, the greatest threat to such innovation and the expression it can yield is a panicked rush to implement precautionary government regulation. Free speech advocate Greg Lukianoff testified to Congress that "the most chilling threat that the government poses in the context of emerging AI is regulatory overreach that limits its potential as a tool for contributing to human knowledge."²⁷

It should come as no surprise then that governments, industries, and experts have focused significant attention on how to regulate AI. Indeed, the sheer number of new regulatory and legislative efforts presents a challenge to meaningful and timely analysis.²⁸ Yet policymakers

should understand the commonalities and trends within recent regulatory efforts and how such efforts not only are confused as to what harms they seek to address but will also hinder free expression.

Legislative Approaches

In the United States, there have been many new bills and working groups focused on the issue of AI, but two bipartisan efforts stand out as emblematic of the government’s attempt to regulate it. The approach taken by Sen. Josh Hawley (R-MO) and Sen. Richard Blumenthal’s (D-CT) AI framework would establish a licensing regime enforceable by the government, make AI companies liable for various harms, and add a long list of legal requirements for how AI must be audited and developed, including the implementation of “safety brakes” and the avoidance of “particularly adverse decisions.”²⁹ This model rejects innovation in favor of government control in its attempt to prevent various nebulous harms to civil rights, children, privacy, election interference, national security, and so on.³⁰ However, there are also more flexible approaches, such as Sen. Mark Warner (D-VA) and Sen. Marsha Blackburn’s (R-TN) Promoting United States Leadership in Standards Act of 2024, which emphasizes the government’s role in supporting the development of technical standards around AI.

As a cautionary tale, the European Union’s (EU) long-standing overregulatory approach has largely crushed European tech innovation. Only 3 of the 50 largest tech companies in the world are based in the EU.³¹ Of the 49 private startups worth more than \$10 billion, only one is from the EU. The EU’s technology sector is held back by general business regulations as well as tech-specific regulations, such as the General Data Protection Regulation, the newly implemented Digital Services Act and Digital Markets Act, and the recently passed EU AI Act. Whether it’s large tech companies, successful startups, or AI-specific companies that are targeted, the EU’s crippling regulations are severely hindering the development of new technologies.³²

Executive Actions

Beyond legislation, there are other ways in which the administrative and regulatory state crafts formal

regulations and exerts informal government pressure to force companies to accept various rules. Most notably, the Biden administration released an executive order that demonstrates many of the fears and concerns over AI. It has been described as an “everything-and-the-kitchen-sink approach” that “represents a significant shift” against innovation with its “red tape wishlist.”³³ By issuing this executive order, the government is attempting to address a host of fears, including those of doomsday AIs, the criminal use of AIs, AIs that can replace or streamline human labor, AIs as products of dominant tech firms, and AIs as threats to particular views of equity, diversity, and justice by way of “algorithmic discrimination.”³⁴

The executive order invokes the Defense Production Act, a law typically meant to ensure that materials essential for war are produced, to restrict AI development by requiring AI developers to follow numerous government standards. The executive order also calls on various agencies, including the Federal Trade Commission and the Department of Justice, to expand their areas of responsibility and authority by creating new rules.³⁵ These administrative agencies also possess significant investigatory and enforcement authority that may be used to prosecute AI developers or coerce them into consent decrees over their supposed harms.³⁶ In sum, the executive order clearly adopts a precautionary stance that believes that “harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks.”³⁷ This is opposed to a pro-innovation view, which acknowledges that AI experimentation should be allowed to advance without requiring substantial government regulation or management since a convincing case has not been made that the new technology is generally harmful. It’s worth noting that the executive order—as well as the even more pessimistic Bletchley Declaration, signed by various nations at the United Kingdom’s AI Safety Summit—is being adopted by significant elements of the technology community, including Google, whose research on Gemini cites both the executive order and the Bletchley Declaration as authoritative in their model safety work.³⁸

Regulatory Takeaways

Based on this cursory review of current government efforts, regulation threatens free expression in two primary ways:

1. By generally making it more difficult for new entrants to join the AI market, regulations will stifle growth and solidify the market dominance of existing AI companies. Many large companies have agreed to rent-seeking regulations that they believe they can manage but that will pose a barrier to new competitors. While policymakers may not intend to strengthen the hold of existing technology companies with such AI regulations, this is the likely effect. The result is that AIs with different or more tolerant viewpoints toward user expression will not be able to take root and emerge, effectively limiting expression through AI to the norms favored by current technology companies, activists, and governments.
2. Beyond the general harms of rent-seeking, AI regulatory efforts are targeting specific types of content they deem to be misinformation and hate speech and demanding that AI development align with specific ideological values and norms that are antithetical to broader expression. This leads to the suppression of AIs that serve users with certain viewpoints. If enacted in the United States, explicit viewpoint regulations could raise constitutional issues, but normative pressure—as we can see in the expectations set by the US executive order and the Bletchley Declaration—may already be having a similar effect. Furthermore, the Brussels effect—that is, the tendency of EU laws to influence the innovation of American companies that want to do business in Europe—is likely limiting the diversity of AI options available to Americans, as AI developers consider international laws with fewer protections for expression when developing a new AI product.

AN INNOVATION-FIRST APPROACH TO AI

Rather than a precautionary, regulate-first approach, the best way to expand expression-supported AI is a risk-based pro-innovation approach to AI development. While some implementations of AI justify extreme caution—such as autonomous military technology with the power to wage war—a risk-based approach acknowledges that most AI applications, especially those involving speech and expression, should be considered innocent until proven guilty.

This concept has been termed “permissionless innovation” and calls for various flexible “soft law” solutions that are appropriate and necessary for AI to flourish.³⁹ These include:

- enforcement of existing law;
- the organic development of social norms and best practices;
- ensuring competition in the marketplace;
- third-party certification; and
- providing Americans with education and AI literacy.

Contrary to popular belief, applying these solutions does not turn AI development into the Wild West, as many laws and regulations still apply. For example, if a criminal uses AI to make his scams more realistic, the scammer can and should still be investigated and prosecuted. In many industries and countries, expansive regulatory and policy systems are already in place, often to the detriment of innovation. In the United States, the National Institute of Standards and Technology, the Food and Drug Administration, the National Highway Traffic Safety Administration, the Consumer Product Safety Commission, the Securities and Exchange Commission, the Federal Trade Commission, the Equal Employment Opportunity Commission, the Federal Aviation Administration, the Department of Justice, and others claim or are exercising regulatory authority over AI products.⁴⁰

A pro-innovation approach also acknowledges that society needs time to learn and adapt to the ways bad actors are abusing new technology. The internet created the opportunity for many new scams that we are now mostly accustomed to (e.g., we usually ignore emails from Nigerian princes). We have learned and will continue to learn how to avoid these bad actors, though some will always succeed, especially as society adapts to new technologies. To this end, the government can make efforts to improve or support education and awareness of new technologies and how they can be abused.

Another soft law that would enable the government to support innovation would be to employ multistakeholder processes that identify best practices and truly voluntary codes of conduct. In a rapidly evolving field such as AI, these best practices would not establish a comprehensive list of what companies can and cannot do—instead, they would focus on outcomes. Relatedly, soft laws can include the collaborative development of norms and principles to guide

AI development. As noted earlier, such efforts to develop norms and identify best practices—especially when backed by government pressure—can result in standards and principles that don't neatly map onto broader liberal values. For example, while a great deal of ethical work on AI has focused on diversity, bias, discrimination, and other potential harms, little appears to be focused on maximizing expression.⁴¹

It is worth mentioning that state and local governments should also adopt pro-innovation policies, as they are increasingly moving to regulate AI without legislative action from Congress.⁴² Unfortunately, local governments could introduce a patchwork of new regulations and restrictions—some that would even conflict with one another—that could chill AI innovation. Therefore, a flexible, pro-innovation approach also requires the federal government to preempt the multiplication of new local laws.

These soft-law approaches will not always get everything right. For example, it could be argued that the norms developed around content moderation are too critical of expression and are playing out poorly in the AI space.

NOTES

1. Marc Andreessen, “Why AI Will Save the World,” a16z, June 6, 2023.
2. Jordi Calvet-Bademunt, “Generative AI Developers Should Commit to Free Speech and Access to Information,” Tech Policy Press, May 7, 2024.
3. “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Proclamation No. 14,110, 88 Fed. Reg. 75,191 (October 30, 2023).
4. “What Is Artificial Intelligence (AI)?,” IBM.
5. Adam Zewe, “Explained: Generative AI,” *MIT News*, Massachusetts Institute of Technology, November 9, 2023.
6. Many experts and organizations have differing ways of categorizing and labeling new AI tools. See, for example, IBM and Adobe, which classify generative AI differently: “Understanding the Different Types of Artificial Intelligence,” IBM, October 12, 2023; and “Generative AI vs. Other Types of AI,” Adobe.
7. Jordi Calvet-Bademunt and Jacob Mchangama, *Freedom of Expression in Generative AI: A Snapshot of Content Policies*

Nevertheless, soft laws are flexible, allowing the market and competition to challenge and correct those mistakes. Especially in the ever-changing and growing AI marketplace, permissionless innovation can support diverse and more expression-focused AI products if given the freedom to develop and operate.

CONCLUSION

AI systems are powerful technologies that can advance human flourishing in a multitude of ways, including enabling new forms of expression. Disruptive new technologies, however, can often cause a moral panic, and AI is no different. To protect Americans' right to free expression, policymakers should not enact precautionary regulations that stifle the development of AI without clear proof of their risk of harm. Furthermore, policymakers should reject efforts to control the ethics and norms around AI-powered expression. Instead, they should favor a robust market of AI tools that can serve as many users and perspectives as possible.

- (Nashville, TN: Future of Free Speech, February 2024).
8. “Top Generative AI Statistics for 2024,” Salesforce, September 2023.
9. Khari Johnson, “People Are Increasingly Worried AI Will Make Daily Life Worse,” *WIRED*, August 31, 2023; and Dan Milmo, Kiran Stacey, and Hibaq Farah, “Nick Clegg Compares AI Clamour to ‘Moral Panic’ in 80s over Video Games,” *The Guardian*, October 31, 2023.
10. Adam Thierer, “Existential Risks and Global Governance Issues around AI and Robotics,” R Street Policy Study no. 291, June 2023.
11. World Economic Forum, *The Global Risks Report 2024*, 19th ed. (Geneva: World Economic Forum, January 2024).
12. Brian Chau, “Google Gemini’s Woke Catechism,” *From the New World* (blog), Alliance for the Future, February 23, 2024; and *Gemini: A Family of Highly Capable Multimodal Models* (Mountain View, CA: Google, December 2023).
13. Amy Kraft, “Microsoft Shuts Down AI Chatbot after It Turned into a Nazi,” *CBS News*, March 25, 2016; and Cade

Metz, “Who Is Making Sure the A.I. Machines Aren’t Racist?,” *New York Times*, updated June 23, 2023.

14. Mike Solana, “Google’s Culture of Fear,” *Pirate Wires*, March 4, 2024.

15. James Bessen, Stephen Michael Impink, and Robert Seamans, “Ethical AI Development: Evidence from AI Startups,” working paper, Center on Regulation and Markets, Brookings Institution, March 17, 2022.

16. “Developing Beneficial AGI Safely and Responsibly,” OpenAI; “Our Principles,” Google AI; “Making Safety Tools Accessible to Everyone,” Meta Llama; and Adam Thierer, “Flexible, Pro-Innovation Governance Strategies for Artificial Intelligence,” R Street Policy Study no. 283, April 2023.

17. “Developing Beneficial AGI Safely and Responsibly,” OpenAI.

18. Jordi Calvet-Bademunt and Jacob Mchangama, *Report: Freedom of Expression in Generative AI—A Snapshot of Content Policies* (Nashville, TN: Future of Free Speech, February 2024).

19. Jacob Mchangama, Abby Fanlo, and Natalie Alkiviadou, *Scope Creep: An Assessment of 8 Social Media Platforms’ Hate Speech Policies* (Nashville, TN: Future of Free Speech, July 2023).

20. David Inserra, “A Guide to Content Moderation for Policymakers,” *Cato Institute Policy Analysis* no. 974, May 21, 2024.

21. *Murthy, Surgeon General, et al. v. Missouri et al.*, No. 23-411, slip op. at 2 (5th Cir. June 26, 2024).

22. Robby Soave, “Homeland Security Agrees That the Disinformation Board Was a Bad Idea,” *Reason*, July 19, 2022.

23. David Inserra, “Lawsuit Alleges More Government Censorship by Proxy—State Department Funds Blacklisting of US Media?,” *Cato at Liberty* (blog), Cato Institute, December 13, 2023.

24. Select Subcommittee on the Weaponization of the Federal Government, *The Weaponization of the National Science Foundation: How NSF Is Funding the Development of Automated Tools to Censor Online Speech “At Scale” and Trying to Cover Up Its Actions* (Washington: House Committee on the Judiciary, February 5, 2024).

25. David Inserra, “A Guide to Content Moderation for Policymakers,” *Cato Institute Policy Analysis* no. 974, May 21, 2024.

26. Nick Vidal, “The AI Renaissance and Why Open Source Matters,” Open Source Initiative, April 11, 2024; and Kevin Roose, “Personalized A.I. Agents Are Here. Is the World Ready for Them?,” *New York Times*, updated November 11, 2023.

27. Jordan Howell, “Greg Lukianoff Testifies before Congress on ‘Chilling Threat’ to Free Speech Posed by AI Regulation,” *Newsdesk*, Foundation for Individual Rights in Education, February 6, 2024.

28. Adam Thierer (@AdamThierer), “Every Lawmaker Gets Their Own Personal AI Bill!,” Medium post, March 1, 2024.

29. Jennifer Johnson, Jayne Ponder, and Conor Kane, “Senators Release Bipartisan Framework for AI Legislation,” *Inside Privacy* (blog), Covington & Burling, September 14, 2023.

30. “Blumenthal & Hawley Announce Bipartisan Framework on Artificial Intelligence Legislation,” Office of Senator Richard Blumenthal, press release, September 8, 2023.

31. “Largest Tech Companies by Market Cap,” Global Ranking, CompaniesMarketCap.

32. Adam Thierer, “Why the Future of AI Will Not Be Invented in Europe,” *Technology Liberation Front* (blog), August 1, 2022; “The EU Artificial Intelligence Act,” EU Artificial Intelligence Act, Future of Life Institute; and Adam Thierer (@AdamThierer), “GDPR & European Innovation Culture: What the Evidence Shows,” Medium post, February 5, 2023.

33. Adam Thierer, “White House Executive Order Threatens to Put AI in a Regulatory Cage,” R Street, October 30, 2023; Jennifer Huddleston, technology policy research fellow, Cato Institute, Testimony before the Subcommittee on Cybersecurity, Information Technology, and Government Innovation of the House of Representatives Committee on Oversight and Accountability, 118th Cong., 2nd sess., March 21, 2024; and Nihal Krishan, “Tech Groups Push Back on Biden AI Executive Order, Raising Concerns That It Could Crush Innovation,” *FedScoop*, November 9, 2023.

34. “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Proclamation No. 14,110, 88 Fed. Reg. 75,191 (October 30, 2023); *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (Washington: White House, October 2022); and Neil Chilson and Adam Thierer, “The Coming Onslaught of ‘Algorithmic Fairness’ Regulations,” Regulatory Transparency Project, Federalist Society, November 2, 2022.

35. “Safe, Secure, and Trustworthy Development and

Use of Artificial Intelligence,” Proclamation No. 14,110, 88 Fed. Reg. 75,191 (October 30, 2023); and Neil Chilson, *Government Overreach in the Presidential Executive Order on Artificial Intelligence*, Testimony before the Subcommittee on Cybersecurity, Information Technology, and Government Innovation of the House of Representatives Committee on Oversight and Accountability, 118th Cong., 2nd sess., March 21, 2024.

36. Neil Chilson and Adam Thierer, “The Coming Onslaught of ‘Algorithmic Fairness’ Regulations,” Regulatory Transparency Project, Federalist Society, November 2, 2022.

37. “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Proclamation No. 14,110, 88 Fed. Reg. 75,191 (October 30, 2023).

38. Brian Chau, “Google Gemini’s Woke Catechism,” *From the New World* (blog), Alliance for the Future, February 23, 2024.

39. Adam Thierer, “Flexible, Pro-Innovation Governance Strategies for Artificial Intelligence,” R Street Policy Study no. 283, April 2023.

40. Adam Thierer, “The Most Important Principle for AI Regulation,” R Street, June 21, 2023.

41. Carlos Ignacio Gutierrez and Gary Marchant, *A Global Perspective of Soft Law Programs for the Governance of Artificial Intelligence* (Phoenix: Arizona State University, 2021).

42. Adam Thierer, “State and Local Meddling Threatens to Undermine the AI Revolution,” *The Hill*, January 21, 2024.



The views expressed in this paper are those of the author(s) and should not be attributed to the Cato Institute, its directors, its Sponsors, or any other person or organization. Nothing in this paper should be construed as an attempt to aid or hinder the passage of any bill before Congress. Copyright © 2024 Cato Institute. This work by the Cato Institute is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.