LAIDLAW SCHOLARS

GEORGETOWN UNIVERSITY
Center for Research & Fellowships

Mapping the AI Ecosystem: Investigating Industry Influence on Federal AI Policy

Rishi Dinesh

Georgetown University

Acknowledgement

I would like to thank the Laidlaw Scholars Foundation and the Georgetown University Laidlaw Scholars Program for supporting my research. I would like to thank Colleen Dougherty for her kindness and guidance. I would like to thank my family for their support, without which I would not be able to do this. Finally, I would like to thank Professor Rajesh Veeraraghavan for his mentorship, wisdom, and encouragement throughout this research process.

1. Introduction

Artificial intelligence (AI) systems are rapidly transforming finance, healthcare, education, defense, and information.[1] While AI may create unprecedented opportunities for growth and prosperity, it also introduces new threats to labor, societal cohesion, and security. To balance these rewards and risks, both advocates and critics of AI have urged the U.S. federal government to adopt AI policy—rules, regulations, and guidelines that govern the development, deployment, and use of AI systems.[2]

Historically, the rapid adoption of general-purpose technologies (e.g., cars, airplanes, social media) has led to cooperation and conflict between industry and government. Current federal AI policy is no different. AI companies use various government relations strategies—public comments, Congressional testimony, campaign contributions, educational meetings, etc.—to advantageously influence the development of AI policy across different federal policymaking entities.

In this paper, I investigate how AI companies shape federal AI policy in the United States. I map the AI ecosystem and identify critical actors, frameworks, and processes which contribute to the formation of AI policy. Using the case study of the Trump administration's AI Action Plan, I find that AI companies successfully advocated for several advantageous provisions.

2. Methods

AI companies deploy various strategies to shape federal AI policy. These strategies differ based on the deployer (i.e., the specific company) and the target (i.e., the specific federal organization). I use two main sources of data to track these efforts. First, I use financial data from OpenSecrets, a research group which compiles information on money in politics, to track

lobbying meetings and expenditures. Second, I analyze responses to various Requests for Information (RFIs) filed by executive agencies. In total, my dataset comprises 36 public comments from AI companies regarding different proposed federal regulations.

My analysis proceeds in three steps. First, I map key elements of the AI ecosystem. Second, I analyze the scale and focus of AI companies' respective lobbying efforts, as well as their respective policy priorities. Finally, I conduct a case study analysis of the Trump administration's AI Action Plan to uncover how companies' policy priorities shaped federal policy.

3. AI Ecosystem

The AI ecosystem includes certain key elements—critical actors, frameworks, and processes which contribute to the formation of AI policy in the United States. Data from RFIs reveals these key elements and their respective influence.

On the industry side, top AI companies (e.g., Amazon, Anthropic, Google, Microsoft, Meta, NVIDIA, OpenAI, Palantir) are key actors who attempt to favorably influence federal regulations. On the government side, executive agencies are key actors who promulgate binding and non-binding regulations and standards in the absence of Congressional action. Technical standards are most commonly set by the National Institute of Standards and Technology (NIST), as well as the Office of Science and Technology Policy (OSTP) and the National Telecommunications and Information Administration (NTIA) to a lesser extent. However, due to the far-reaching implications of AI systems, especially LLMs, agencies like the Department of Energy (DOE) and the U.S. Copyright Office (USCO) have also set out their own guidance.[3]

Several industry-led organizations, frameworks, and standards also play key roles in shaping AI policy. In the absence of federal guidance, AI companies have developed internal

industry-wide practices, creating some standardization in areas such as AI governance, synthetic content verification, and international interoperability.

For example, the Coalition for Content Provenance and Authenticity (C2PA), whose members include Amazon, Microsoft, Meta, Google, and OpenAI, publishes a set of widely adopted industry standards on content verification. Microsoft and Google both explicitly recommend that C2PA standards be adopted as federal rules.

Moreover, standards developed by the Organization for Economic Cooperation and Development (OECD), International Organization for Standardization (ISO), and International Electrotechnical Commission (IEC) influence industry practices and preferences. Almost every top AI company cites ISO/IEC 42001 as a model for domestic standards. The OECD AI Principles, which represent the first published set of intergovernmental principles on AI, are also widely cited in industry arguments on what the future of federal AI policy should look like.[4]

4. Lobbying

Data from OpenSecrets reveals the respective lobbying influence of different AI companies. Since 2020, Meta has spent the most money lobbying the federal government ($90.9 million), followed respectively by Amazon ($84.5 million), Google ($58.2 million), Microsoft ($44.2 million), Palantir ($17.3 million), OpenAI ($2.6 million), NVIDIA ($2.2 million), and Anthropic ($1.4 million).[5]

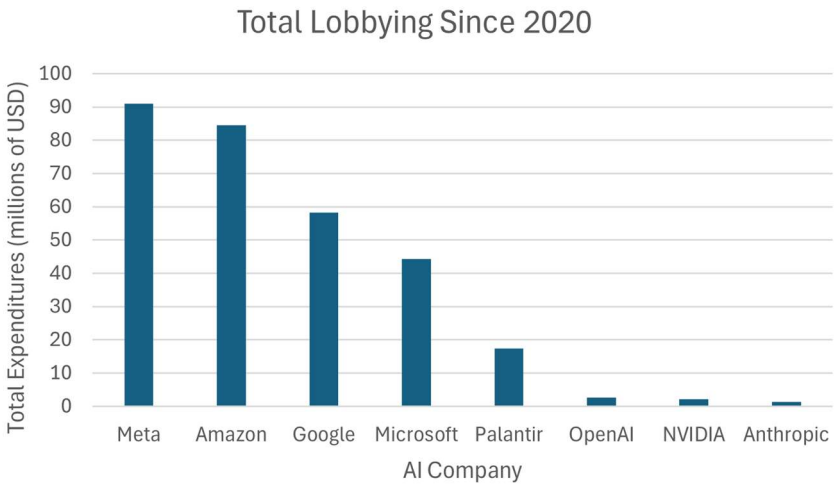

In 2024, these companies disclosed Congressional lobbying activity on an aggregate of 401 bills or resolutions. However, these sums represent total lobbying activity. To measure "AI-related" lobbying activity, I filter the OpenSecrets data for bills which mention artificial intelligence, AI, algorithms, or deepfakes in the title. Of the aforementioned 401 bills and resolutions, 84 (20.9%) were "AI-related." On those bills and resolutions, companies filed an aggregate of 287 reports under the Lobbying Disclosure Act (LDA).

	All Bills	AI-related bills	AI-related reports/issues
Amazon	108	7	24
Anthropic	4	4	8
Google	103	22	60
Meta	43	11	38
Microsoft	126	30	124
NVIDIA	1	1	2
OpenAI	9	9	31
Palantir	7	0	0
Total	401	84	287

Of the top AI companies, Microsoft lobbied for the most AI-related bills and resolutions (30), followed by Google (22), Meta (11), and OpenAI (9). Microsoft also filed the most AI-related LDA reports (124), followed by Google (60), Meta (38), and OpenAI (31). Notably, while every other top AI company filed between 2–3.44 reports per AI-related legislation, Microsoft displayed a much higher rate of 4.13.

The filtered data also reveals each company's policy priorities for 2024. OpenAI, Anthropic, and NVIDIA lobbied exclusively on AI-related legislation. Roughly one fourth of Microsoft, Meta, and Google's lobbying efforts focused on AI-related legislation. Less than a tenth of Amazon's lobbying was AI-related, and Palantir did not lobby for any AI-related legislation—its lobbying focused on defense and foreign policy bills. These differences reflect market positioning. OpenAI and Anthropic derive revenue almost exclusively from AI products.[6] Google, Meta, Microsoft, and Amazon, on the other hand, are primarily known for products in related fields (e.g., web browsing, social media, systems software, retail).

5. AI Action Plan

In February 2025, the National Science Foundation opened an RFI on the development of an AI Action Plan. Amazon, Anthropic, Google, Meta, Microsoft, Palantir, and OpenAI each submitted a public comment detailing their policy priorities and vision for the AI Action Plan. Five months later, the Trump administration published the finalized Plan, which featured 30 policy goals organized under three pillars—innovation, infrastructure, and international leadership.

To determine the extent to which AI companies shaped the AI Action Plan, I measure how closely the AI Action Plan reflects companies' policy priorities. These similarities should not be attributed to coincidence. Under the Administrative Procedure Act, executive agencies are

required to consider all relevant, timely-submitted comments, which demonstrates that industry arguments are read and factored into final regulations.[7]

Pillar I—Remove Red Tape and Onerous Regulation

The Action Plan takes aim at local, state, and federal laws which are "unduly restrictive to innovation." It accomplishes this by reducing federal funding to states with "burdensome" AI regulatory regimes, functionally enacting federal preemption of state laws.[8] This focus on preempting cumbersome and conflicting regulations closely reflects the complaints of several AI companies. Microsoft argues that regulatory fragmentation diverts resources from innovation to compliance. Amazon similarly argues that businesses may struggle to navigate diverging standards, which they identify as a "substantial risk." Both Amazon and Google explicitly advocate for federal preemption of what they describe as a "patchwork of [state-level] laws." Furthermore, Microsoft and Google point out that regulatory fragmentation even in related areas such as cybersecurity and privacy may slow progress on AI.[9]

The most descriptive vision for federal preemption appears in OpenAI's response. Their ideal system would create a "single, efficient 'front door' to the federal government" for AI companies. OpenAI advocates for the federal government to support industry in two ways. First, federal agencies should coordinate expertise to provide guidance on classified intelligence, government contracting, and model evaluations. Second, the federal government should preempt state regulations on frontier model security—OpenAI explicitly identifies California SB 1047 as an example.[10]

In discussing federal preemption, many AI companies advocate against state regulations which are inflexible and binding. OpenAI articulates their ideal vision as a "voluntary partnership" between the federal government and private companies. Microsoft argues for

voluntary, industry-led standards, which they argue would bolster definitional clarity and predictability.[11] Both of these suggestions align with the Action Plan's focus on flexibility and deregulation.

Pillar I—Build an AI Evaluations Ecosystem

The Action Plan seeks to "publish guidelines and resources through NIST at DOC, including CAISI, for Federal agencies to conduct their own evaluations of AI systems." It also advocates that NIST, along with the DOE and NSF, should support the development of the "science of measuring and evaluating AI models."[12] This provision responds to requests from Amazon and Anthropic for support on model evaluation capabilities. The similarities between the comments and the Action Plan are striking. For example, Amazon advocates for federal support to "advance the science behind capability evaluations, for national security risks."[13]

Anthropic also articulates this focus on national security. They state (in bold font) that "the federal government must develop robust capabilities to rapidly assess any powerful AI system, foreign or domestic, for potential national security uses and misuses."[14] In an earlier provision (*Ensure that Frontier AI Protects Free Speech and American Values*), the AI Action Plan reflects these concerns: "Led by DOC… publish evaluations of frontier models from the People's Republic of China."[15] Thus, the Action Plan's language and framing on model evaluations closely mirrors industry priorities.

Pillar I—Encourage Open-Source and Open-Weight AI

The Action Plan aims to create a "supportive environment for open models," which represents the first major federal endorsement of open-source and open-weight models. Of the top AI companies, only Meta advocates for prioritizing open-source models. They point out that

"Meta has been at the forefront of open source for over a decade," emphasizing the benefits to national security, efficiency, and innovation.[16]

Meta's top priority in its RFI response is supporting the adoption of open models. In the introduction itself, Meta devotes an entire paragraph to explaining the importance of open-source development and Meta's role in that space. Its arguments on open-source models represent the first and largest section of their comment. Given that no other top AI company even discussed this issue, it is likely that the AI Action Plan's provisions on open-source and open-weight AI were heavily influenced by Meta.

Pillar II—Create Streamlined Permitting for Data Centers, Semiconductor Manufacturing Facilities, and Energy Infrastructure while Guaranteeing Security

The Action Plan takes aim at burdensome, resource-intensive regulations and permitting processes which "make it almost impossible to build [AI] infrastructure in the United States."[17] Specifically, the Plan establishes new exemptions under the National Environmental Policy Act (NEPA) for certain data center-related actions; applies the FAST-41 streamlined permitting process to data center and data center energy projects; streamlines and reduces environmental regulations; and explores a potential expedited nationwide data center permit under Section 404 of the Clean Water Act (CWA).

These clauses clearly mirror industry arguments from a wide range of companies. Amazon, Anthropic, Google, Meta, Microsoft, and OpenAI all point to the need for streamlined permitting for data center and energy infrastructure.[18] In fact, many of the Action Plan's specific policy tools borrow from companies' suggestions. OpenAI urges the federal government to create "categorical exclusions to the National Environmental Policy Act."[19] Meta's two policy recommendations in this area both appear in the final Action Plan—a new "Nationwide Permit

for Data Center Uses" and a streamlined permitting process under CWA Section 404.[20] The extreme specificity of these incorporations suggests that the federal government intentionally derived parts of the AI Action Plan from companies' public comments.

Pillar III—Strengthen AI Compute Export Control Enforcement

The Action Plan seeks to reduce circumvention of export controls by creating new tools for detection and coordinating with the intelligence community.[21] This provision is especially interesting because Google, OpenAI, and Meta all advocate *against* it. Google argues that export controls designed under the Biden administration's Diffusion Rule were unduly restrictive and onerous.[22] OpenAI proposes a modification to the Rule which would "expand the number of countries in Tier I," substantially reducing the breadth of export controls.[23] And Meta advocates for exempting open-source models from export controls. Instead of adopting either of these suggestions, the Trump administration prioritized national security arguments.[24] This example demonstrates the limits of industry influence—while AI companies' arguments are often incorporated or factored into federal policy, scaling back national security tools remains contentious.

6. Conclusion

Given the unprecedented and transformative consequences of AI, it is more important than ever to understand the drivers of AI policy. This paper makes several findings to enhance that understanding. AI policy in the United States is driven by interactions between key federal science agencies and top AI companies. AI companies strategically use several lobbying tactics to insert their priorities into policymaking processes in Congress and the executive branch. AI companies are not monolithic in their tactics or their preferences.

All of these findings are evident within the Trump administration's AI Action Plan. The Plan heavily reflected industry preferences in a number of consequential domains. It also even adopted some specific policy recommendations made by AI companies, demonstrating the impact of industry influence on federal AI policy. However, industry priorities are sometimes overridden by other interests such as national security. Ultimately, more research is needed to uncover the specific channels and methods AI companies use to influence federal policy. By doing so, researchers, policymakers, and the public may better understand how to shape the future of such a consequential technology.

---

[1] Martin Neil Baily and Aidan T. Kane, "Harnessing AI for Economic Growth: Insights from Electricity, Finance, Health Care, and Information Sectors," Brookings Institution, April 8, 2025, https://www.brookings.edu/articles/harnessing-ai-for-economic-growth/.

[2] Lenovo, *AI Policy Essentials: Frameworks for Innovation*, Lenovo Glossary, https://www.lenovo.com/us/en/glossary/ai-policy.

[3] "AI and Copyright," August 30, 2023, U.S. Copyright Office, https://www.regulations.gov/document/COLC-2023-0006-0001/comment.

[4] Microsoft, "MSFT Response to NTIA RFC on AI Accountability," 2023: 12, https://www.regulations.gov/comment/NTIA-2023-0005-1337. Microsoft, "Response of Microsoft Corporation to NIST RFI on an Artificial Intelligence Risk Management Framework," 2021: 7–9, https://www.nist.gov/system/files/documents/2021/09/16/ai-rmf-rfi-0088.pdf.

[5] "OpenSecrets," OpenSecrets, accessed August 14, 2025, https://www.opensecrets.org/.

[6] Roald Larsen, "How Anthropic Makes Money: The Business Model Explained." *Untaylored*, July 21, 2024. Zachary DeWitt, "Notorious: OpenAI's Revenue Breakdown (A Closer Look)," *Notorious*, July 18, 2024.

[7] Administrative Conference of the United States, *IIB-014: Notice-and-Comment Rulemaking* (Bulletin No. 014), May 2021, https://www.acus.gov/sites/default/files/documents/IIB014-Rulemaking.pdf.

[8] "America's AI Action Plan," 2025, The White House: 3, https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.

[9] Microsoft, "Comments Received in Response to: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 10, https://files.nitrd.gov/90-fr-9088/Microsoft-AI-RFI-2025.pdf. Amazon, "Comments of Amazon: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 6, https://files.nitrd.gov/90-fr-9088/Amazon-AI-RFI-2025.pdf.

[10] OpenAI, "Comments Received in Response to: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 6, https://files.nitrd.gov/90-fr-9088/OpenAI-RFI-2025.pdf.

[11] Microsoft, "Artificial Intelligence (AI) Action Plan," 13.

[12] "America's AI Action Plan," The White House, 10.

[13] Amazon, "Artificial Intelligence (AI) Action Plan," 6.

[14] Anthropic, "Comments Received in Response To: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 3–4, https://files.nitrd.gov/90-fr-9088/Anthropic-AI-RFI-2025.pdf.

[15] "America's AI Action Plan," The White House, 4.

[16] Meta, "Comments Received in Response to: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 1–8, https://files.nitrd.gov/90-fr-9088/Meta-AI-RFI-2025.pdf.

[17] "America's AI Action Plan," The White House, 14.
[18] Amazon, "Artificial Intelligence (AI) Action Plan," 4. Anthropic, "Artificial Intelligence (AI) Action Plan," 7. Google, "Comments Received in Response To: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 2025: 3, https://files.nitrd.gov/90-fr-9088/Google-RFI-2025.pdf. Meta, "Artificial Intelligence (AI) Action Plan," 11. Microsoft, "Artificial Intelligence (AI) Action Plan," 3. OpenAI, "Artificial Intelligence (AI) Action Plan," 12.
[19] OpenAI, "Artificial Intelligence (AI) Action Plan," 13.
[20] Meta, "Artificial Intelligence (AI) Action Plan," 11.
[21] "America's AI Action Plan," The White House, 21.
[22] Google, "Artificial Intelligence (AI) Action Plan," 3–4.
[23] OpenAI, "Artificial Intelligence (AI) Action Plan," 7–9.
[24] Nury Turkel, "AI, National Security, and the Global Technology Race: How US Export Controls Define the Future of Innovation," *Hudson Institute*, March 24, 2025, https://www.hudson.org/national-security-defense/ai-national-security-global-technology-race-how-us-export-controls-define-nury-turkel.