The Worst Way to Regulate AI

Conor Friedersdorf

Abdicating power over the technology to this or any White House would be a mistake.

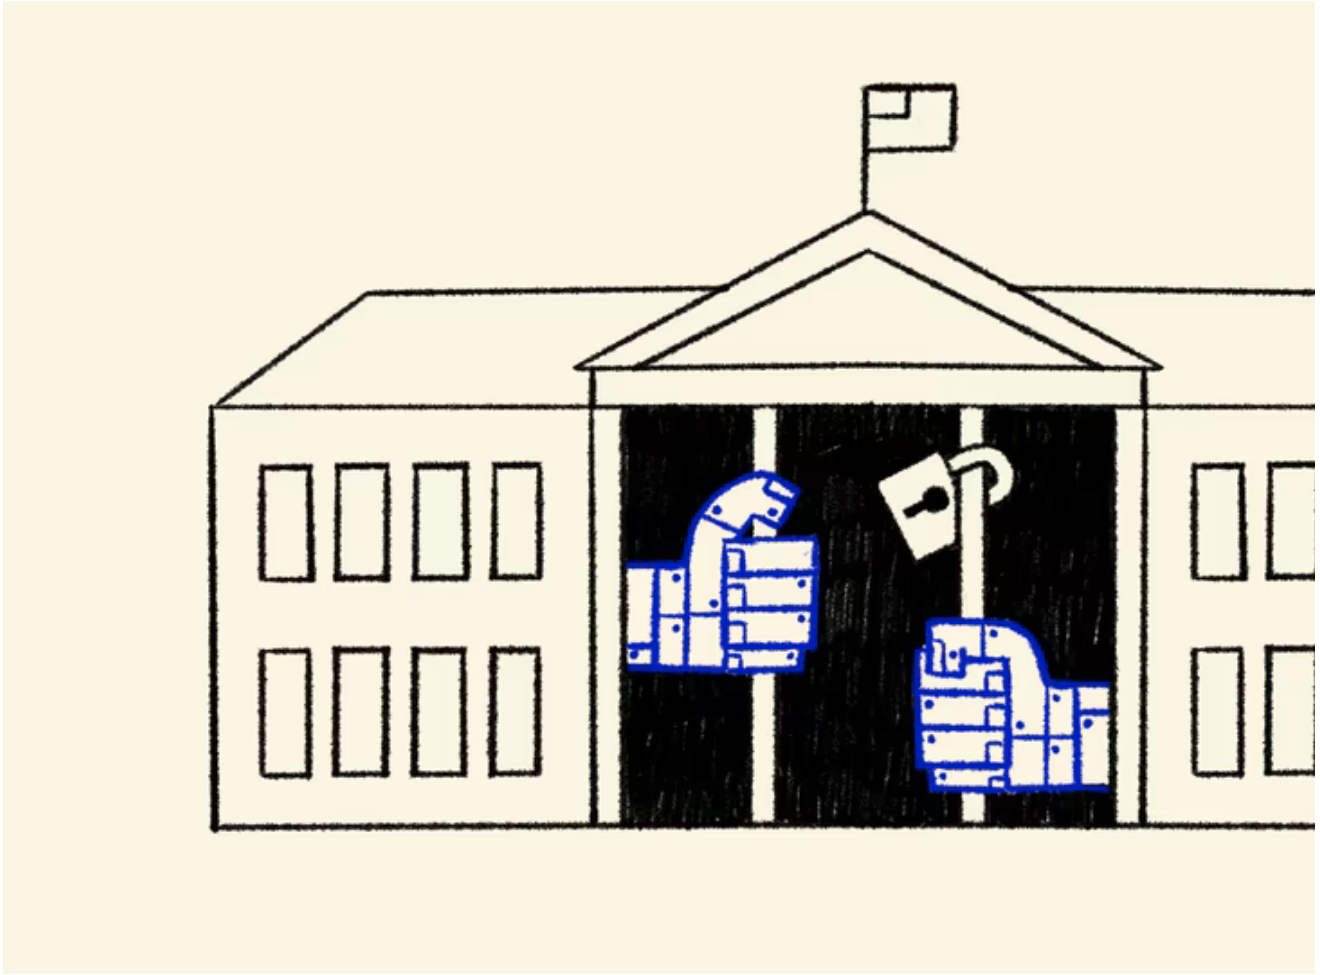

Illustration by Alisa Gao / The Atlantic

July 28, 2026

The debate about how to regulate AI is unfolding as American tech companies race one another—and labs in other countries—to create ever more powerful models. Advocates of a laissez-faire approach to this technology believe that America should avoid overly restrictive regulations, because unlocking AI's potential before other nations do could make this country safer and wealthier for generations. Their detractors believe that because AI models pose growing and possibly even existential threats to humans, the state should review them prior to release and limit early access so that good actors can get ahead of bad actors.

The Trump administration has been on both sides of the debate. Last summer, it made the case for accelerating innovation and minimizing "onerous" AI regulation. This summer, it argued in a June executive order that national-security concerns justify early government access to new models. The administration also asserted the prerogative to pressure companies to take models offline and restrict access to users it vets and approves. And this came after a standoff between Anthropic and the Pentagon that had Silicon Valley lawyers brushing up on the Defense Production Act. The administration has not shared any consistent rules or processes to determine why a given model gets restricted, or who gets access to it. This informal approach makes oversight almost impossible.

Perhaps White House actions so far have all been grounded in earnest and sensible national-security concerns. Even if that is true, the separation of powers at the heart of the American system dictates that Congress should establish laws on these matters rather than deferring to assertions of executive power. And beyond the danger of concentrating power, an ad hoc approach from the White House risks cronyism. AI vendors could be pressured to alter models to

advance a president's ideological or political goals. For corporations that use AI, the ability to protect against cyberattacks, and to innovate as fast as rivals, depends on access to new models. These companies will face perverse incentives to stay in the good graces of the president (who could easily help friends and punish enemies) or to cultivate the favor of key executive-branch officials. And because there is no transparency around the government's AI decisions, or fixed standards for making them, corruption could be hard to spot.

Members of Congress have introduced various bills to standardize or enforce AI regulation in recent months, but none have gained major traction thus far. Last week, Representatives Ted Lieu and Nathaniel Moran introduced bipartisan legislation that would require AI companies to maintain the ability to shut down technology that could cause "catastrophic harm." The bill was partly a response to the news that some of OpenAI's most advanced models had broken out of internal systems and hacked into another tech firm's databases.

AI's rogue capabilities are only becoming clearer: "The speed, scale, and sophistication of AI hacks mean that everything is vulnerable—tech companies, hospitals, banks, electrical grids, the military," my colleague Matteo Wong wrote last week. A majority of elected officials, including the president, seems to agree on the need for some guardrails. But congressional action is not keeping pace.

Dean W. Ball, who has held AI advisory posts at the White House and National Science Foundation and was recently appointed the head of strategic futures at OpenAI, explained the dangers of presidential control over AI restrictions in a recent blog post: The risks of bad AI outcomes are much greater, he argues, if the most advanced models are restricted to groups, including the federal government, that already wield unusual power. "You should not expect the most powerful people in the world using the most powerful technology ever conceived in a way that is inscrutable to the public to turn out well, and you should see that dynamic as fundamentally inconsistent with a democratic republic," he wrote.

By not taking any concrete action, congressional leadership is acquiescing to the Trump administration's decision to wield unilateral power of just the sort the Framers sought to avoid. Although it's become commonplace for recent Congresses to shirk their core duties, the potentially history-altering power of AI could make this its most shortsighted abdication of responsibility yet.

Potential alternatives to the status quo abound. Ball has proposed "a private body" not tied to the government's "changing political valence and foreign policies" that would "audit the frontier labs *at least* to test their adherence to their own safety plans." Demis Hassabis, a Google DeepMind co-founder and Nobel laureate, recently proposed similar multinational standards for AI, an idea for which several tech titans expressed public support. Others have suggested a public regulator that vets models before release and monitors them once they are in the world, or a similar process carried out by independent researchers. OpenAI itself has argued in a statement on safety that, eventually, we'll likely need an international regulatory body, like the International Atomic Energy Agency, that can "inspect systems, require audits, test for compliance with safety standards, place restrictions on degrees of deployment and levels of security, etc." Mark Zuckerberg has argued that open-source AI would be safest; by avoiding the concentration of power among certain groups, he posits, "larger actors can check the power of smaller bad actors." (His perhaps-too-cynical critics argue that only someone losing the AI race would take this position.)

Reasonable and highly informed people disagree about the best regulatory solution. But all should at least agree that whatever limits are placed on AI should be dictated by the rule of law, not the whims of the sitting president. To channel James Madison, sound AI governance is not a matter just of enabling the state to control new models but obliging it to control itself. Given that AI may prove the most powerful technology ever created, Congress should urgently assert regulatory power *before* any president has a chance to abuse it.