

How Do You Regulate an Industry You Don't Understand?

Columbia researchers say we can govern AI even as we're still learning how it works.

By

David J. Craig

|

Oct. 19, 2026



David Plunkert

Americans' attitudes toward artificial intelligence are deeply conflicted. We appreciate that AI enables us to work more efficiently, but worry about it eliminating

millions of jobs. We love the lightning-fast photo edits, but not the data centers driving up our electric bills. We believe that the technology may one day help scientists tackle cancer and climate change, but only if it doesn't wipe out humanity with a novel bioweapon first.

Every new technology brings both opportunity and peril, of course. And it is the job of policymakers to craft regulations that maximize the former while limiting the latter. But how do you regulate technology that is so complex and rapidly evolving that even its creators don't fully understand it? This is the dilemma now faced by elected leaders across the US. With constituents demanding that AI companies be held accountable for the safety and societal impacts of their products, and the leaders of those same companies insisting that only they are capable of monitoring their systems' performance, lawmakers are left wondering where to start.

Fortunately, scholars have been anticipating this moment for years, studying how the AI industry might be regulated in ways that avoid stifling innovation while ensuring that AI tools serve the public good. Columbia researchers have published a number of books and papers on the topic, exploring how new laws could push AI companies to invest more resources in safety and transparency, how existing laws might be applied to mitigate AI's new risks, and how businesses of all types can maintain customers' trust by voluntarily adopting responsible AI-use policies.

[Steven Cohen](#), a SIPA professor who studies public policy and management, says that the first step in effectively regulating any powerful new industry is to give independent experts a window into its inner workings, no matter how difficult that is. The US government has created entirely new agencies, like the Federal Aviation Commission and the Food and Drug Administration, to perform such watchdog functions in the past. Cohen believes that the risks posed by artificial intelligence — which range from sophisticated cybercrime and bioterrorism attacks to deepfake scams and AI-dependency issues — warrant the creation of a similar new agency. "It's clear that AI is going to affect nearly every aspect of our lives, so the public has a right to know what kinds of models are being developed, and what their benefits and risks could be," he says. The federal government is best equipped to do that monitoring, Cohen's work suggests. "It will be expensive to do right, a major undertaking," he says. "The government will need to build the technical expertise and auditing capacity to evaluate these systems and keep pace as they evolve. Some scientists who helped to develop AI will probably have to be lured away from the industry to take part."

In recent weeks, many observers, including the editors of the *New York Times*, have advocated for establishing a federal agency to oversee the AI industry. Few anticipate this happening anytime soon, though. President Donald J. Trump, upon returning to office last year, scuttled plans that Joseph Biden had set in motion in 2023 to create federal safety standards for AI and to require government-overseen testing of models before release; Trump administration officials have since stated that the nation's frontier AI firms should be trusted to police themselves. Still, public opinion in the US has swung sharply in favor of regulating the industry of late, especially since several of the firms, including OpenAI, Anthropic, Google DeepMind, and Meta, have acknowledged that powerful new agents they've been developing have gone on hacking sprees. The ensuing debate about the unprecedented dangers posed by these novel AI agents (*Could they launch nuclear missiles? Or disable our electric grid?*), combined with growing concerns about AI-related job losses and the construction of thousands of sprawling, electricity-devouring data centers across the country, has pushed majorities of Americans to now say they distrust the big AI companies and expect that their products will ultimately do more harm than good. "The political backlash that's built up against the industry is grassroots and bipartisan, which tells me it's likely to persist," says Cohen. If the federal government doesn't crack down on the companies soon, he predicts, non-profit organizations may step in to try to fill the gap, perhaps analyzing publicly available information about AI-related incidents and issuing their own warnings about emerging risks. "A well-financed NGO could have a significant impact in this way," he says. "Even if its assessments weren't legally binding, a watchdog group could influence public opinion and the AI companies' business prospects."

And how exactly would independent observers evaluate AI systems, if mandated by law? Would they need to peer over the shoulders of the AI companies' scientists as they train and test their models, inspecting every miscue and coding fix, and then issuing licenses prior to release? Would it suffice to monitor the performance of AI products after they're released, documenting adverse events and conducting post-mortem forensics to prevent them from recurring? Or might a company's performance be measured indirectly, perhaps by analyzing their investments in safety and quality-control initiatives? [Neel Guha](#), a Columbia law professor who studies the AI industry, sees potential for all of these strategies. For now, though, he favors a moderate approach centered on mandated adverse-event reporting, under which AI companies would be legally bound to report any instances of harm caused by their products to a newly created independent body that would have broad

authority to investigate. “Our most urgent need is to find out how often frontier models are malfunctioning and precisely what’s happening in those instances,” Guha says. “Mandated reporting is the fastest, most cost-effective way to get that information out into the public.”

Guha says that if such an auditing body existed today, either as part of a government agency or a private entity reporting to one, we’d know much more about the recent hacking incidents committed by the frontier labs’ AI agents. “The information that the companies have publicly released so far is lacking in many crucial details,” he says, “which goes to show that a purely voluntary reporting system isn’t enough.”

Most AI experts interviewed by *Columbia Magazine* say that while this year’s hacking incidents are alarming, they don’t prove that the technology is evolving beyond human control. In fact, many scholars argue that if the big AI firms simply invested more resources in evaluating and verifying the safety of their AI models, rather than speeding ahead blindly, they could control their systems. They are suspicious, therefore, of calls from AI executives for a government-forced pause on development, which could freeze out competition from smaller companies and solidify the big firms’ market dominance. “Right now, we’re at a stage where we still need a clearer understanding of the dangers posed by these technologies and what the causes are,” says Guha, who is the co-author of one widely cited study that scrutinizes a number of oft-proposed measures for regulating the industry. “Our initial steps at regulation should be aimed at gathering that information.” He says that more aggressive regulations could have unintended consequences, including monopolization.

Scholars say that a number of laws already on the books could be used to rein in the activities of AI firms. [Lina Khan](#), a Columbia law professor who chaired the Federal Trade Commission from 2021 to 2025, has argued that AI companies and their executives may face extensive civil liability if their products harm others, for example. “Much of the commentary from the executives suggests they currently operate in a legal vacuum, as if AI has wiped clean the regulatory map,” Khan wrote in a [recent *New York Times* op-ed](#). “In fact, their business practices — including how they develop and deploy their technologies — are already subject to longstanding laws: Product liability laws apply when goods are defective or dangerous, while tort law establishes basic duties for companies to not act negligently or recklessly. Consumer protection laws prohibiting unfair or deceptive practices can also kick in

when businesses injure others by failing to take adequate precautions.” She argues that prosecutors could even hold AI executives criminally responsible if their systems engage in egregiously illegal or dangerous activities, with state AGs potentially taking the lead.

The public’s fascination with rogue AI agents, and the existential threat they may pose, can obscure a number of more immediate risks associated with AI, according to Columbia experts. Evidence suggests that the technology is already contributing to job losses, cybercrime, and youth mental-health challenges, for example. And scholars have been exploring ways to mitigate these harms. [Joseph Stiglitz](#), a Nobel Prize-winning Columbia economist, has written about how the nation’s tax code might be updated to funnel a sizeable portion of the AI industry’s profits, which are expected to be enormous, into wage subsidies and other social welfare initiatives. [Claire Huntington](#) ’96LAW, a Columbia law professor, has studied how legislators could draw on family law, which defines when human relationships become abusive or exploitive, to set standards for how AI systems must treat the humans who use them. “Many people are now entering into what are essentially relationships with their chatbots, whether for friendship or therapy or even romantic or sexual connections,” Huntington says, “and this can make them vulnerable in all sorts of ways.”

She recommends that lawmakers require AI companies to make their chatbots less addictive, especially for children and adolescents, and to remind users that the tools are merely simulating emotional intimacy. “A common misconception is that personal relationships are entirely private and beyond the reach of the law, but that’s just not true,” Huntington says. “Family law regulates relationships in which there are power imbalances and potential for abuse, as between adults and children, and it gives us concrete strategies for how we might regulate these AI companies.”

Ultimately, the decisions that individual banks, hospitals, insurance companies, news organizations, and other businesses make about how and when to deploy AI will determine the shape of our economy moving forward, points out [Daniel Dobrygowski](#), an adjunct professor at SIPA and who leads the cybersecurity and AI legal team at the network security firm Sophos. That means that in the absence of a major regulatory push in Washington, DC, the safeguards that companies voluntarily put in place for AI in the coming months and years will have far-reaching consequences. In his recent book, [Technology Governance](#), Dobrygowski provides practical advice for business leaders considering how to integrate AI systems into their workflows. The

most important tip: Remember what your customers expect in terms of data security and privacy, and hold those goals paramount even while looking for ways to increase productivity with AI. “You need to remain attentive to the basics,” he says. “Your business fundamentals shouldn’t change just because you’re using AI.”

Regardless of how strictly the US government regulates artificial intelligence, Dobrygowski predicts, the winners and losers in an AI-driven economy will increasingly be determined by the level of trust that ordinary people have in different companies to protect their data and in other ways use the technology responsibly. “One of the big myths in corporate culture is that you have to wait for regulations to tell you what kinds of safety guardrails to adopt,” he says. “But that’s not true. Business leaders today know what their customers’ concerns are around AI. They can act responsibly. They make good decisions on their own.”

Read more from **David J. Craig**



[Guide to school abbreviations](#)

[All categories >](#)

Read more from **David J. Craig**