

What the Aviation Industry Can Teach Us About Regulating Artificial Intelligence

Joshua M. Javits*

When a technology such as AI potentially poses serious risks to the public, safety cannot completely depend on an industry's willingness to police itself. As conservative economics icon Milton Friedman has framed the issue, "Government is essential...for determining the 'rules of the game' and as an umpire to interpret and enforce the rules." The aviation industry offers important lessons for how we can oversee artificial intelligence to help ensure public safety.

Understandably, Washington, DC, should not regulate AI exactly as it regulates airlines. AI is fundamentally different from aviation, with risks ranging from privacy and discrimination to cybersecurity, economic disruption, biosecurity and warfare. The more useful axiom is that aviation demonstrates how government and private industry can jointly oversee AI and share responsibility for ensuring public safety.

Should AI be regulated at all?

AI industry leaders came to the Trump Administration asking for some regulation, based on their recognition of the extreme risks AI potentially poses to public safety and the deficiencies of self-regulation to mitigate those risks. The real question is not whether AI should be regulated but rather what government should regulate and what it should leave to the private sector to innovate and develop commercially. Unfortunately, industry leaders came away from their meeting with the Administration sorely disappointed. The knee-jerk response from the Administration was to reject any regulation at all, leaving entirely to the private sector the untenable task of mitigating the technology's enormous risks to the public.

In the airline industry it is the risks and outcomes, rather than the precise technology, that are regulated. The lessons from aviation regulation for AI are that government should establish safety standards and ensure these standards are being met, even as the industry develops further applications of the technology.

The government establishes the safety floor.

The Federal Aviation Administration does not simply ask airlines to promise that they will operate safely. Before an airline can begin operations, it must demonstrate compliance with federal safety requirements and possess the ability to detect and manage hazards. After certification, government oversight continues through surveillance, inspections and other safety processes.

The aviation industry—both carriers and manufacturers—is responsible for most of the engineering and operational work necessary for air safety. But government retains the authority

to determine whether minimum safety requirements are being met and to intervene when they are not.

Government oversight matters, because some AI risks extend far beyond an individual company's customers. In biosecurity, for example, AI could lower barriers to accessing information and knowledge relevant to biological weapons. In warfare, AI could influence intelligence, targeting or other decisions involving the use of force, including its initiation. Such risks involve public health and national security, not simply commercial interests, and no single company can fully contain them. Government cannot eliminate AI risks. Importantly, however, it can establish minimum safety requirements for systems whose failure could have extraordinary consequences.

The appropriate model is shared responsibility.

AI developers should be accountable for understanding, testing and documenting the capabilities and limitations of the systems they build. Companies deploying AI in consequential settings should also be accountable for determining whether the technology is appropriate for its intended use. Government oversight, too, must take an accountable approach. Regulators should have clearly defined responsibilities, transparent standards and the authority necessary to enforce them. Accountability is therefore critical for both the companies developing and deploying AI and the government agencies deemed responsible for protecting the public.

This is also where voluntary self-regulation has limits. Internal safety teams, independent audits and board oversight can be valuable tools to assess and mitigate risks. But what happens when a company does not comply? In aviation, a regulator has legal authority to require corrective action and, in appropriate circumstances, restrict an operator's ability to continue operating. Under a voluntary AI system, that ultimate enforcement mechanism not only is less clear but may not even exist. Consequently, self-regulation can be only one layer of AI safety.

The regulatory framework should also be proportional to risk.

An AI system that writes thank-you notes should not face the same regulatory requirements as one used in a hospital, power grid, financial system, biological research facility or military decision-making process. The principle should be simple: The consequences of misuses or failures of AI should determine the intensity of oversight.

As in aviation, certification cannot be the end of the process.

Aircraft remain subject to oversight by the Federal Aviation Administration after entering service because conditions change and new hazards emerge. Manufacturers must support their products with advisories to the carriers that use their products as well as amendments to operation and maintenance manuals based on experience and feedback. AI presents an even stronger case for continuing oversight, because models can be retrained, connected to new data, given new capabilities or incorporated into different products.

AI needs an institutional feedback loop: test before deployment, monitor after deployment, investigate serious failures, require corrective action, learn from experience and update standards.

The National Transportation Safety Board is another key part of the aviation model. When an accident occurs, the NTSB investigates and issues safety recommendations designed to prevent similar accidents. The purpose is not merely to assign blame; it is to require remedies that make the entire system safer. Although the FAA is not obliged to follow NTSB recommendations, the NTSB adds a layer of recognized expert independent safety analysis.

Finally, the NASA's Aviation Safety Reporting System is a voluntary, confidential, non-punitive, safety reporting system that receives safety reports from pilots, air traffic controllers, dispatchers, cabin crew, maintenance technicians and drone operators. Its reports describe unsafe occurrences, near misses or close calls, hazardous situations and descriptions of lessons learned and best practices. It provides a level of immunity to the whistle blowers to enable the reporting.

AI needs an equivalent learning mechanism. When an AI system causes serious harm, a structured process should exist for determining what happened, whether the failure was foreseeable, whether safeguards were adequate and whether the incident should change standards for other systems. The goal is not to make failure impossible; it is to make failures visible, learnable and less likely to recur.

The National Institute of Standards and Technology's AI Risk Management Framework reflects much of this philosophy—organizing risk management around governing, mapping, measuring and managing risks throughout an AI system's life cycle. Its underlying concept is important: Safety is an ongoing responsibility, not a one-time promise.

Thus the White House agreement with the major AI technology companies can be useful without being sufficient. Voluntary commitments can encourage better corporate practices, independent evaluation can expose problems and corporate boards can demand accountability. But when consequences can extend to biological security, critical infrastructure, national security and armed conflict, government must play a role in the safety regime.

The history of aviation suggests how government oversight of AI can work.

Government need not design AI, operate AI companies or dictate every engineering decision. It can establish minimum safety requirements, require evidence that high-risk systems meet them, maintain continuing oversight, investigate serious failures and intervene when necessary. What is absolutely essential is an enforcement mechanism similar to the regime imposed under the Federal Aviation Act, in which a government authority can impose fines, suspensions or revocations of operational authority on both individuals and companies if standards are violated.

Private innovation and public oversight can coexist. AI is unlikely to follow the same path as aviation; its applications are broader, its development is faster and its risks are more diffuse. The question is not whether government can guarantee that AI will never cause harm. It cannot. The

question is whether we are willing to build a system in which companies innovate, independent experts test, regulators oversee and both developers and government are accountable for safety.

Companies can and should be part of the solution, but they should not be the only line of defense when failure can affect the public. The aviation industry model offers a useful starting point: test before deployment, assign responsibility, establish accountability, monitor continuously, investigate failures, learn from experience, scale oversight according to risk and give government the authority—and the responsibility—to act when public safety is at stake.

A partnership between AI developers and government to effect social responsibility will enable this amazing and fast-advancing technology to flourish while affording some assurances of risk mitigation.

* Joshua M. Javits is a neutral mediator and arbitrator. He is the immediate past president of the National Academy of Arbitrators. He was previously chair and member of the National Mediation Board, grievance chair for the International Monetary Fund and president of the National Association of Railroad Referees.