

## Comments of the International Center for Law & Economics

*Office of Science and Technology Policy RFI, Regulatory Reform on Artificial Intelligence OSTP-TECH-2025-0067, 90 FR 46422*

October 27, 2025

### Authored by:

**Eric Fruits** (Senior Scholar, International Center for Law & Economics)

**Ben Sperry** (Senior Scholar, International Center for Law & Economics)

**Kristian Stout** (Director, Innovation Policy, International Center for Law & Economics)

# Comments of the International Center for Law & Economics

*Eric Fruits, Ben Sperry, Kristian Stout\**

## Executive Summary

These comments from the International Center for Law & Economics analyze the economic effects of artificial-intelligence (AI) governance and recommend a policy of strategic forbearance. Because AI is a general-purpose technology with a great degree of uncertainty about its potential future pathways, it would be preferable to apply existing, technology-neutral statutes—such as those governing fraud, discrimination, and product safety—while gathering empirical evidence, rather than to enact premature, prescriptive mandates. This approach preserves the “option value” of waiting for better information before imposing irreversible regulatory costs.

Current barriers to AI deployment stem from two primary sources: federal regulatory mismatch and state-level market fragmentation. Many existing federal regulations are structurally incompatible with AI systems because they embed human-centric assumptions. In transportation, Federal Aviation Administration (FAA) rules for pilot certification and National Highway Traffic Safety Administration (NHTSA) standards for vehicle controls assume a human operator, preventing the approval of autonomous systems. In health care, the Food and Drug Administration’s (FDA) premarket approval process is designed for static physical devices, which conflicts with continuously learning diagnostic algorithms. Similar mismatches exist in financial services under the Equal Credit Opportunity Act and in health-care reimbursement under Medicare’s telehealth rules.

At the state level, a growing patchwork of conflicting AI regulations—such as those enacted in Colorado, California, and New York City—creates substantial economic burdens with high compliance costs. This fragmentation also creates extraterritorial effects, where one state’s restrictive law may become a *de facto* national standard, burdening interstate commerce.

This paper proposes a two-part framework to address these barriers.

1. **Federal administrative action:** Agencies should use existing administrative flexibility—including waivers, pilot programs, and conditional approvals—to permit AI deployment and data gathering. The Commercial Space Launch Amendments Act of 2004 provides a successful precedent for this model. Agencies must also conduct a systematic audit to identify and modernize human-centric regulations, replacing them with performance-based standards.

---

\* Eric Fruits and Ben Sperry are senior scholars with the International Center for Law & Economics (ICLE). Kristian Stout is ICLE’s director of innovation policy. ICLE has received financial support from numerous companies, foundations, and individuals, including firms with interests both supportive of and in opposition to the ideas expressed in this and other ICLE-supported works. Unless otherwise noted, all ICLE support is in the form of unrestricted, general support. The ideas expressed here are the authors’ own and do not necessarily reflect the views of ICLE’s advisors, affiliates, or supporters.

2. **Targeted congressional preemption:** Congress should enact legislation that preempts state and local laws regulating the design, development, training, and validation of AI models. This preemption, modeled on the Airline Deregulation Act, would harmonize the national market. It would simultaneously preserve states' authority to enforce general, technology-neutral laws (e.g., consumer-protection statutes) and the authority of federal agencies like the Federal Trade Commission (FTC) and the Equal Employment Opportunity Commission (EEOC) to police demonstrated harms.

## I. Introduction and Overview

The Office of Science and Technology Policy (OSTP) is seeking public input to identify federal laws, regulations, and administrative processes that unnecessarily restrict the development and adoption of artificial intelligence (AI) development in the United States.<sup>1</sup> The International Center for Law & Economics (ICLE) is a nonprofit, nonpartisan research organization whose core mission is to promote the application of law & economics methodologies to inform public-policy discussion. Our work focuses on developing intellectually rigorous, data-driven analyses to foster efficient policy solutions that enhance consumer welfare and global economic growth. ICLE previously submitted comments to OSTP regarding national priorities for AI<sup>2</sup> and the development of an AI action plan,<sup>3</sup> and to the National Telecommunications and Information Administration (NTIA) regarding AI accountability policies.<sup>4</sup>

OSTP's request for information (RFI) asks respondents to identify federal statutes, regulations, and policies that unnecessarily hinder AI development, deployment, and adoption. The request organizes potential barriers into five categories: regulatory mismatch, structural incompatibility, lack of regulatory clarity, direct hindrance, and organizational factors.

Our analysis rests on three economic propositions. First, as was the case with the introduction of wired electric power or the commercial internet, AI is a general-purpose technology whose full economic and social applications remain uncertain.<sup>5</sup> Promulgating technology-specific regulations in such circumstances may foster what economists call targeting errors: rules that are simultaneously too broad (prohibiting beneficial uses) and too narrow (failing to address actual harms).<sup>6</sup>

Second, the compliance costs that arise from fragmented or overly prescriptive regulation function as barriers to entry, particularly for smaller firms that cannot easily absorb the fixed costs of legal review and documentation.

Third, existing legal frameworks already address most AI-related harms through technology-neutral statutes. Indeed, agencies like the Federal Trade Commission (FTC), Equal Employment

---

<sup>1</sup> *Request for Information: Regulatory Reform on Artificial Intelligence*, OFF. OF SCI. & TECH. POLICY (OSTP-TECH-2025-0067, 90 FR 46422, Sep. 26, 2025), [hereafter "RFI"].

<sup>2</sup> *Response of the International Center for Law & Economics, Request for Information, National Priorities for Artificial Intelligence*, INT'L. CTR. LAW & ECON. (OSTP-TECH-2023-0007, Jul. 7, 2023) available at <https://laweconcenter.org/wp-content/uploads/2023/07/OSTP-AI-Comments.pdf> [hereafter "Priorities Comments"].

<sup>3</sup> *Comments of the International Center for Law & Economics: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan*, INT'L. CTR. LAW & ECON. (Mar. 14, 2025) available at <https://laweconcenter.org/wp-content/uploads/2025/03/OSTP-AI-2025-comments-v-1.pdf> [hereafter "Action Plan Comments"].

<sup>4</sup> *Response of the International Center for Law & Economics to the AI Accountability Policy Request for Comment: National Telecommunications and Information Administration*, INT'L. CTR. LAW & ECON. (Docket No. 230407-0093, Jun. 12, 2023) available at <https://laweconcenter.org/wp-content/uploads/2023/06/NTIA-AI-Comments-final.pdf> [hereafter "Accountability Comments"].

<sup>5</sup> Accountability Comments 17.

<sup>6</sup> Priorities Comments 5.

Opportunity Commission (EEOC), and Food and Drug Administration (FDA) already enforce existing laws against practices such as deception, discrimination, and unsafe products, regardless of whether these issues arise from algorithmic tools or traditional methods.

These principles support several concrete recommendations:

- The federal government should conduct a systematic audit of regulations that embed human-centric assumptions incompatible with adaptive or autonomous systems, prioritizing modernization in transportation, health care, and financial services.
- Agencies should expand their use of existing administrative flexibilities—waivers, pilot programs, conditional approvals, and time-limited experimental authorities—to enable lawful AI deployment, while gathering evidence about actual risks and benefits.
- The Office of Management and Budget (OMB) should coordinate interagency guidance to clarify how existing statutes apply to AI systems, with the aim of preventing duplicative or conflicting interpretations across agencies.
- Congress should establish targeted preemption of conflicting state laws that regulate the design, development, and training of AI models in order to prevent a compliance patchwork that fragments the national market.
- Agencies should incorporate formal innovation-impact assessments into cost-benefit analysis for AI-related rulemakings, quantifying effects on market entry, competition, and U.S. technological competitiveness.
- Future AI regulations should include sunset provisions and mandatory periodic review to ensure rules adapt as the technology evolves.

## II. Strategic Forbearance and the Option Value of Information

A core principle of the law & economics methodology holds that regulation should (1) correct demonstrated market failures; (2) match the regulatory tool to the nature of the market failure; and (3) do so only when intervention benefits exceed costs.<sup>7</sup>

Judge Richard Posner, formerly of the 7th U.S. Circuit Court of Appeals, and Cass Sunstein of Harvard Law School have each written extensively on the problem of regulating under uncertainty.<sup>8</sup>

---

<sup>7</sup> STEPHEN BREYER, *REGULATION AND ITS REFORM* (1982) 191 (“Our examination of market defects, classical modes of regulation, and alternative regimes suggest that regulatory failure sometimes means a failure to correctly match the tool to the problem at hand.”), 184 (“It should be painfully apparent that whatever problems one has with an unregulated status quo, the regulatory alternatives will also prove difficult.”)

<sup>8</sup> Richard A. Posner, *Antitrust in the New Economy*, 68 *ANTITRUST L. J.* 925 (2001), (“The real problem lies on the institutional side: the enforcement agencies and the courts do not have adequate technical resources, and do not move fast enough, to cope effectively with a very complex business sector that changes very rapidly.”); Isaac Ehrlich & Richard A. Posner, *An Economic Analysis of Legal Rulemaking*, 3 *J. LEGAL STUD.* 257, 277 (1974), (“An important cost of legal regulation by means of rules is thus the cost of altering rules to keep pace with economic and technological change.”); Cass R. Sunstein, *Maximin*, 37 *YALE J. REG.* 940 (Jul. 29, 2020), (“Observers (including regulators) are in circumstances of Knightian uncertainty, where they cannot assign probabilities to imaginable outcomes.”).

When regulators possess limited information about a new technology's risks and benefits, prescriptive rules often impose costs that exceed their welfare gains.

AI presents an acute example of this uncertainty problem. AI technologies encompass everything from simple pattern recognition in manufacturing quality control to complex language models that generate text and code, as well as advanced multimodal systems capable of producing realistic images, artwork, and dynamically generated video content. Its applications span health-care diagnostics, financial underwriting, transportation logistics, agricultural monitoring, and scientific research. No single regulatory framework can efficiently govern such heterogeneity.

Adam Thierer of the R Street Institute has documented how premature and overly rigid regulation of nascent technologies can create path dependencies that persist long after the original rules become obsolete. His analysis of telecommunications regulation in the 1980s and 1990s illustrates the ways that prescriptive mandates delayed the deployment of digital networks and mobile telephony.<sup>9</sup>

The economic concept of *option value* helps to explain why regulatory forbearance can enhance consumer welfare. When regulators act prematurely, they foreclose the option to gather more information through market experimentation. By waiting, they preserve the ability to craft more precise rules based on observed outcomes, rather than speculation about hypothetical risks.<sup>10</sup> As ICLE's Gus Hurwitz and Geoffrey Manne explain:

Regulation as a discovery process, in its simplest formulation, asks regulators to consider that they might be wrong. That they might be asking the wrong questions, collecting the wrong information, analyzing it the wrong way—or even that Congress has given them the wrong authority or misunderstood the problem that Congress has tasked them to address. And, in response to these concerns, regulation as a discovery process asks regulators to build them into the regulatory process itself. This is because regulation does not operate like a competitive market. If Amazon Prime had not been the successful idea that it was, consumers would not have adopted it, other firms would have obtained competitive advantage over Amazon, and the idea would have fallen into the dustbin of history. But when an agency promulgates a rule pursuant to APA processes, that rule takes on the force of law—good or bad, there is no market mechanism by which it will succeed [or] fail.<sup>11</sup>

---

<sup>9</sup> See, e.g., Adam Thierer, *The Future of Innovation: Is This the End of Permissionless Innovation?*, DISCOURSE (Jan. 6, 2021), <https://www.discoursemagazine.com/p/the-future-of-innovation-is-this-the-end-of-permissionless-innovation>.

<sup>10</sup> See, e.g., Kenneth J. Arrow & Anthony C. Fisher, *Environmental Preservation, Uncertainty, and Irreversibility*, 88 Q. J. ECON. 312 (1974), (explaining that, when a decision is irreversible and future benefits and costs are uncertain, immediate action eliminates the opportunity to learn more before committing resources; by waiting, society retains an “option value,” i.e., the value of preserving flexibility to act later when uncertainty has been reduced; delaying regulation (forbearance) can be welfare enhancing if immediate action would lock in a costly or suboptimal rule and prevent learning from future information or market responses).

<sup>11</sup> Justin (Gus) Hurwitz & Geoffrey A. Manne, *Regulation as a Discovery Process* (Oct. 16, 2024), <https://ssrn.com/abstract=4721112>.

The Commercial Space Launch Amendments Act of 2004 (CSLAA) provides a successful precedent for the regulation-as-a-discovery-process approach.<sup>12</sup> Congress recognized that applying traditional aerospace regulations, which were designed for government-operated vehicles, would prevent the development of commercial human spaceflight. The CSLAA set safety guidelines for U.S. commercial human spaceflight under the Federal Aviation Administration (FAA), requiring FAA oversight at all launches and landings but prohibiting safety regulations until 2023, unless a serious incident occurred.

Under CLSAA's rules, space-tourism operators must inform participants of launch and reentry risks in writing, disclose vehicle safety records, and obtain participants' informed consent. This was not deregulation, as the FAA retained authority to license launches, investigate accidents, and enforce safety standards for the general public. Operators remain liable for damages under common law tort. The CSLAA simply delayed prescriptive regulation of specific design features until there were enough flights to provide an empirical basis for requirements.

Two decades later, this forbearance policy has enabled the growth of a commercial space industry that includes reusable rockets, private space stations, and lunar landers—technologies that did not exist when the CSLAA passed.<sup>13</sup> The FAA now possesses sufficient flight data and operational experience to develop performance-based safety standards grounded in evidence, rather than speculation.

The CSLAA helped to establish a thriving U.S. commercial space industry by granting investors and innovators a stable and predictable regulatory environment. The law's demonstrated success may now offer lessons for AI governance. Agencies overseeing AI-related markets should similarly monitor real-world deployments: observing how developers test systems, how deployers update models, and what failures occur in practice. Regulators can use these insights to craft targeted rules that address actual problems without limiting beneficial uses.

The CSLAA also demonstrates that the common critique of regulatory forbearance—that it amounts to “doing nothing” in the face of risk—is misplaced. Strategic forbearance is not synonymous with regulatory inaction or abdication. It is an active governance strategy focused on the gathering of information. The CSLAA *sequenced* the regulation rather than *eliminating* it. It created a legal framework for learning, data collection, and supervised experimentation. This reframes the policy debate from an either-or proposition of whether to regulate into a framework for developing evidence-based regulation under conditions of high uncertainty.

The cost of premature regulation is not speculative. In 2022, the Information Technology & Innovation Foundation calculated that, over a period of a decade, a patchwork of state privacy laws imposes costs exceeding \$1 trillion on firms that operate in multiple states, with small businesses

---

<sup>12</sup> Pub. L. No. 108-492, 118 Stat. 3974 (2004).

<sup>13</sup> See, e.g., *The Space Report 2025 Q2 Highlights Record \$613 Billion Global Space Economy for 2024, Driven by Strong Commercial Sector Growth*, SPACE FOUND. (Jul. 22, 2025), <https://www.spacefoundation.org/2025/07/22/the-space-report-2025-q2> (the global space economy reached \$613 billion in 2024, and the commercial sector made up roughly 78% of that total).

bearing at least \$200 billion of this total.<sup>14</sup> The report documented how divergent definitions of “personal information,” “consent,” and “automated decision-making” force companies to maintain separate compliance systems for each jurisdiction.

### **III. Regulatory Mismatch and Modernization Priorities**

The RFI identifies regulatory mismatch as occurring when existing requirements are based on assumptions about human operation that do not align with AI capabilities or operational models.<sup>15</sup> When regulations embed human-centric assumptions, they create not just paperwork burdens but genuine resource costs, as firms are forced to either forgo beneficial technologies or maintain parallel compliance systems.

#### **A. Transportation: Human-Operator Assumptions**

The FAA’s pilot-certification regime illustrates an example of the problem of regulatory mismatch. Title 14 of the Code of Federal Regulations, Part 61, establishes requirements for airman certificates. Section 61.3 mandates that no person may act as a required pilot-flight crewmember unless that person holds the appropriate certificate, has it in their physical possession, and presents it for inspection upon request. Section 61.23 further requires an FAA-issued medical certificate. Section 61.39 specifies prerequisites for practical tests, including ground training in topics such as the physiological factors that may affect pilot performance.

While these requirements are perfectly sensible for human pilots, they are absurd for autonomous flight systems. An AI does not have a cardiovascular system and cannot be impaired by alcohol or prescription drugs.

The National Highway Traffic Safety Administration (NHTSA) faces an analogous problem with autonomous vehicles. The existing Federal Motor Vehicle Safety Standards assume human drivers. Standard No. 114, for example, requires theft-protection devices, including a key-locking system and audible warnings.<sup>16</sup> Standard No. 135 requires light-vehicle brake systems to be activated by foot control.<sup>17</sup>

NHTSA has responded to the emergence of autonomous driving systems by allowing manufacturers to file petitions under 49 USC § 30113, which permits limited exemptions for vehicles that provide an equivalent or superior level of safety. But the petition process is resource intensive. For example, General Motors submitted a petition in January 2018 seeking exemptions for a self-driving vehicle

---

<sup>14</sup> Daniel Castro, Luke Dascoli, & Gillian Diebold, *The Looming Cost of a Patchwork of State Privacy Laws*, INFO. TECH. & INNOVATION FOUND. (Jan. 2022), available at <https://www2.itif.org/2022-state-privacy-laws.pdf>.

<sup>15</sup> RFI 46423 (“Most existing Federal regulatory regimes and policy mechanisms were developed before the rise of modern AI technologies. As a result, they often rest on assumptions about human-operated systems that are not appropriate for AI-enabled or AI-augmented systems.”)

<sup>16</sup> 49 CFR § 571.114.

<sup>17</sup> 49 CFR § 571.135.

design based on its Chevrolet Bolt platform.<sup>18</sup> The agency spent 15 months reviewing the GM petition before seeking public comment. The agency did not issue a final decision before GM withdrew the petition in 2020. Earlier this year, U.S. Transportation Secretary Sean Duffy confirmed that the exemption process has been “bogging developers down in unnecessary red tape that makes it impossible to keep pace with the latest technologies.”<sup>19</sup>

The economic solution is not to eliminate safety requirements but to refocus them on performance outcomes, rather than specific human-centered processes. Instead of requiring that a vehicle have a steering wheel operated by a licensed driver, regulations could require that the vehicle demonstrates the capability to maintain lane position, avoid obstacles, and respond to traffic-control devices.

## **B. Health Care: Static Device Assumptions and Continuous Learning**

The FDA’s regulatory framework for medical devices presents a different form of mismatch. The premarket approval process, codified in 21 CFR Part 814, was designed for physical devices like pacemakers and artificial joints. A manufacturer demonstrates safety and effectiveness through clinical trials of a fixed device design, receives approval, and then produces identical copies. Any significant modification to the approved device requires a new submission for FDA review.

This model creates structural problems for AI-enabled medical software that learns from real-world data. Consider a diagnostic algorithm that analyzes radiology images to detect potential tumors. The system’s accuracy improves as it processes more cases and incorporates feedback from radiologist reviews. Under traditional device regulation, these improvements could constitute modifications that require FDA review, effectively freezing the algorithm at its initial, less-accurate state.

The FDA has attempted to address this problem through its Predetermined Change Control Plan guidance, finalized in December 2024 and updated in August 2025.<sup>20</sup> The guidance allows manufacturers to specify in advance what types of modifications they plan to make post-market, the protocols to implement and validate those changes, and the monitoring plan to ensure continued safety and effectiveness.

A more systematic solution would acknowledge that AI-enabled medical devices are fundamentally different from static physical devices. Rather than requiring upfront approval of a fixed system, the FDA could approve a manufacturer’s testing, monitoring, and validation infrastructure. If the manufacturer demonstrates robust protocols for detecting degraded performance, validating updates

---

<sup>18</sup> *Cruise, GM to Seek U.S. Okay for Self-Driving Vehicle Without Pedal, Steering Wheel*, REUTERS (Oct. 21, 2020), <https://www.reuters.com/business/autos-transportation/cruise-gm-seek-us-okay-self-driving-vehicle-without-pedal-steering-wheel-2020-10-21>.

<sup>19</sup> David Shepardson, *NHTSA Streamlining Self-Driving Car Exemption Reviews*, INS. J. (June 16, 2025), <https://www.insurancejournal.com/news/national/2025/06/16/827924.htm>.

<sup>20</sup> *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions: Guidance for Industry and Food and Drug Administration Staff*, FOOD & DRUG ADMIN. (Aug. 18, 2025), <https://www.fda.gov/media/166704/download>.

against clinical benchmarks, and monitoring real-world outcomes, the FDA could permit continuous improvement within defined boundaries.

### C. Financial Services: Model Opacity and Explanation Requirements

The Equal Credit Opportunity Act (ECOA), implemented through Regulation B, requires creditors who deny credit or take other adverse action to provide specific reasons for that action.<sup>21</sup> The Consumer Financial Protection Bureau's (CFPB) Circular 2022-03 clarifies that this requirement applies to creditors using complex algorithms, rejecting arguments that model opacity could excuse the lack of specific reasons.<sup>22</sup>

This creates a regulatory mismatch when lenders rely on machine-learning models that are not easily interpretable. Deep neural networks and similar algorithms do not produce transparent “reasons” for their decisions—unlike models that yield clear decision trees, regression coefficients, or ranked variable importance.<sup>23</sup> As a result, lenders face costly choices: they can forgo these non-interpretable models even when such models better distinguish creditworthy from uncreditworthy applicants; they could construct *post-hoc* explanations after decisions are made; or they could invest in developing model-agnostic explanatory tools. Yet none of these strategies directly advance the underlying policy goal: ensuring that credit-approval rates for protected classes do not significantly differ from those for similarly qualified applicants in comparison groups.

## IV. State Regulatory Fragmentation and the Economics of Preemption

While federal regulatory mismatch and structural incompatibility create barriers to AI deployment, the proliferation of inconsistent state laws poses a distinct economic threat.<sup>24</sup> A patchwork of state requirements imposes costs that grow multiplicatively with the number of jurisdictions. The National Conference of State Legislatures (NCSL) reports that state lawmakers introduced more than 1,100 AI-related bills nationwide in the 2025 legislative sessions.<sup>25</sup> Dozens of states have

---

<sup>21</sup> 12 CFR Part 1002.

<sup>22</sup> *Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms*, Consumer Financial Protection Circular 2022-03, CONSUMER FIN. PROT. BUREAU (May 26, 2022), (“ECOA and Regulation B do not permit creditors to use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions”)

<sup>23</sup> See, e.g., Rabia Saleem, Bo Yuan, Fatih Kurugollu, Ashiq Anjum, & Lu Liu, *Explaining Deep Neural Networks: A Survey on the Global Interpretation Methods*, 513 NEUROCOMPUTING 165 (2022).

<sup>24</sup> Action Plan Comments 3 (“This definitional challenge is further complicated by the emerging patchwork of state and local AI regulations that has emerged in the absence of federal guidance. Developers and deployers of AI systems who operate across jurisdictional boundaries face substantial compliance challenges as a result of this regulatory fragmentation. The proliferation of potentially conflicting state regulations also creates significant legal uncertainty that disproportionately burdens smaller innovators and startups, as these entities often lack the resources to navigate complex regulatory environments. This has the potential to further entrench the market positions of larger incumbents.”)

<sup>25</sup> *Artificial Intelligence 2025 Legislation*, NAT'L CONF. STATE LEGS. (Jul. 10, 2025), <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation>.

enacted legislation or established task forces to study AI regulation. Just four—California, Illinois, New York, and New Jersey—have introduced more than 25% of all state-level AI-related bills.

The definitional inconsistency across these laws is economically significant. Colorado’s SB 24-205, which was amended to take effect June 30, 2026,<sup>26</sup> requires a developer of a high-risk AI system to take reasonable care to protect consumers from any known or reasonably foreseeable risks of “algorithmic discrimination.” The law defines “high-risk artificial intelligence system” as any AI system that makes or is a substantial factor in making a “consequential decision,” which it defined as decisions that have legally material or similarly significant effects on the provision of education, employment, financial services, government, health care, housing, insurance, or legal services.

California began enforcing new rules earlier this year that regulate how employers use AI and automated decisionmaking systems in hiring and employment.<sup>27</sup> These regulations, part of the state’s Fair Employment and Housing Act, aim to prevent discrimination by ensuring that AI tools don’t unfairly disadvantage job applicants or employees based on protected traits like race, gender, or age. Employers must keep records of their AI-related employment data for four years and are responsible for any discriminatory outcomes, even if the tools come from outside vendors.

Similarly, New York City’s Local Law 144 imposes bias-audit requirements specifically for automated employment-decision tools.<sup>28</sup> The law defines such tools as computational processes, derived from machine learning, statistical modeling, data analytics, or AI, that issue simplified outputs that are a substantial factor in employment decisions.

These divergent definitions force developers to maintain separate compliance systems for different jurisdictions. An employment-software provider serving clients in New York City, Colorado, and California must track three different sets of definitions, documentation requirements, audit procedures, and disclosure obligations. As more states continue to explore legislation in this area, this problem of fragmentation will only compound.

### **A. The Extraterritoriality Problem and Constitutional Economics**

State AI regulations do not simply impose local costs on local firms. Because AI models are typically developed for national or global deployment, a restrictive state law effectively sets the standard for all states. In practice, this extraterritorial application violates the principles of federalism and creates economic inefficiency.

---

<sup>26</sup> Michael J. Laszlo, *Colorado’s AI Law Delayed Until June 2026: What the Latest Setback Means for Businesses*, CLARK HILL (Aug. 28, 2025), <https://www.clarkhill.com/news-events/news/colorados-ai-law-delayed-until-june-2026-what-the-latest-setback-means-for-businesses>.

<sup>27</sup> Danielle Ochs, *10 FAQs About California’s New Algorithmic Discrimination Rules*, OGLETREE DEAKINS (Sep. 16, 2025), <https://ogletree.com/insights-resources/blog-posts/10-faqs-about-californias-new-algorithmic-discrimination-rules>.

<sup>28</sup> New York City, Local Law No. 144 of 2021, N.Y.C. Admin. Code § 20-870 *et seq.* (2021), <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>

For its part, California has taken a piecemeal approach across multiple statutes. AB 2013 requires generative-AI providers to publish documentation about the data used to train their models.<sup>29</sup> SB 942 requires persons creating, coding, or otherwise producing generative-AI systems to provide a publicly accessible detection tool and to make certain disclosures for content produced with the technology.<sup>30</sup> AB 853 requires large online platforms to provide consumers with an easy and conspicuous way to determine if content includes provenance information indicating that it was created or altered by an AI system.<sup>31</sup>

These rules can have significant consequences outside of California. Consider AB 2013, which requires generative AI providers to publish documentation about the data used to train their models. The statute applies to any provider that makes a generative-AI system available to California residents. This requirement is effectively a *de facto* national standard because it is technically infeasible to train different models for different states.

The dormant Commerce Clause prohibits states from enacting laws that unduly burden interstate commerce, even absent federal legislation.<sup>32</sup> The doctrine rests on an economic rationale: when states impose costs on out-of-state economic activity, they externalize those costs, while capturing local benefits, leading to socially excessive regulation. In *Pike v. Bruce Church, Inc.*, the U.S. Supreme Court held that a state regulation is invalid if “the burden imposed on [interstate] commerce is clearly excessive in relation to the putative local benefits.”<sup>33</sup>

While the Court did recently narrow the scope of the dormant Commerce Clause, the majority in the fractured decision *National Pork Producers Council v. Ross* decision<sup>34</sup> also upheld the *Pike* balancing test, which allows challenges to nondiscriminatory state laws where the burdens on interstate commerce outweigh the putative benefits to the state. The majority did not think, however, that changes in market share resulting from some companies choosing to exit the market rather than comply with the law was sufficient to establish a substantial burden on commerce. In other words, so long as out-of-state entities retained the ability to choose whether to serve the particular market, there was no dormant Commerce Clause problem.

In the AI context, this balance weighs heavily against AI-specific state regulations. For developers, avoiding states with burdensome AI laws may not even be possible. For instance, New York’s RAISE Act applies to models that are “deployed” in New York, even if the base model was created

---

<sup>29</sup> Cal. Assemb. Bill 2013 (Irwin) (2024) (to be codified at Cal. Civ. Code §§ 1798), available at <https://sjud.senate.ca.gov/system/files/2024-06/ab-2013-irwin-sjud-analysis.pdf>.

<sup>30</sup> Cal. S.B. 942, 2023-24 Sess. (Cal. 2024), <https://legiscan.com/CA/text/SB942/id/3021807>.

<sup>31</sup> Cal. Assemb. Bill 853 (Wicks) (2025) (Ch. 674, Statutes of 2025) (to be codified at Cal. Bus. & Prof. Code §§ 22757.3.1 et seq.), <https://legiscan.com/CA/text/AB853/id/3245321>.

<sup>32</sup> See Barry Friedman & Daniel T. Deacon, *A Course Unbroken: The Constitutional Legitimacy of the Dormant Commerce Clause*, 97 VA. L. REV. 1877 (2011).

<sup>33</sup> 397 U.S. 137 (1970).

<sup>34</sup> 598 U.S. 356 (2023).

elsewhere.<sup>35</sup> Moreover, open-source AI models are often used by downstream users in ways unforeseeable by the original developers; in fact, that is their entire point. There is no way these developers could avoid being subject to state laws on AI if entities within states with such laws use their base model.

Unlike the pork producers in *Ross*, who could segregate their operations to serve the rest of the country while avoiding California, AI developers will be subject to the most stringent laws regardless of their intent to serve that market. Rather than a choice of whether to serve a particular state, developers would effectively have to choose between complying with the most burdensome state regulations or not releasing their models to the public at all.

By contrast, the local benefits of such rules are likely to be marginal. It would be exceedingly difficult for new AI-specific regulations to generate safety or consumer-protection benefits beyond those already offered by generally applicable, technology-neutral rules. Costly administrative procedures like safety protocols, disclosures, and audits have little correlation to safety.

Nonetheless, as covered in the next section, the limitations of the dormant Commerce Clause should encourage Congress to enact a statute preempting state laws that would effectively destroy a common national market.

## **B. A Targeted Preemption Framework**

Congress should enact targeted preemption of state laws that regulate the design, development, training, and validation of AI models.<sup>36</sup> This preemption should be express—stated clearly in statutory text—and should be limited to avoid unnecessary displacement of state authority.

At a high level, such preemption should cover regulations that target the development of AI, *per se*, whether explicitly or implicitly. In other words, preemptive federal legislation should prohibit states from imposing requirements on the core processes involved in creating AI and AI-enabled products (as opposed to how users may *employ* AI technologies). Thus, federal preemption should cover subjects such as the data used to train AI models, the algorithms or techniques used in model development, and the testing and validation procedures applied during development.

States should, however, continue to enforce general laws against fraud, deception, discrimination, and safety violations using their existing authorities under consumer protection, civil rights, and tort law. Moreover, the federal statute should authorize federal agencies with sector-specific expertise to establish performance-based standards for AI systems in their domains.

---

<sup>35</sup> New York 6453–A (Mar. 5, 2025).

<sup>36</sup> Action Plan Comments 4 (“[T]he AI Action Plan should establish clear federal guidelines that preempt contradictory state and local regulations, while setting minimum transparency standards appropriate to each category of AI application. The goal of such standards should be to protect consumers without imposing excessive compliance burdens that might stifle innovation. The aim should be to foster functional markets where customers can access the services they demand, not to initiate a new cottage industry for AI compliance lawyers.”)

The Airline Deregulation Act of 1978 provides a successful template. Section 41713(b)(1) of Title 49 prohibits states from enacting or enforcing laws “related to a price, route, or service of an air carrier.”<sup>37</sup> The Supreme Court has interpreted this language broadly to prevent states from using indirect regulation to affect airline operations.

Congress should adopt similarly broad preemption language for AI, prohibiting state laws related to the design, development, training, or validation of AI systems. This would ensure that developers face uniform national standards, rather than fragmented state requirements.

## **V. Technological Neutrality and the Sufficiency of Existing Law**

The RFI asks whether existing policy frameworks are appropriate for AI applications. Most AI-related harms are already addressable through technology-neutral statutes governing fraud, discrimination, unfair commercial practices, and product safety.<sup>38</sup>

### **A. FTC Authority Over Deceptive and Unfair Practices**

Section 5 of the Federal Trade Commission Act prohibits “unfair or deceptive acts or practices in or affecting commerce.”<sup>39</sup> The FTC has made clear that it believes Section 5 authority extends to AI systems. For example, earlier this year, the commission finalized an order against DoNotPay, a company marketing itself as the world’s first “robot lawyer.”<sup>40</sup> The complaint alleged that DoNotPay deceptively claimed its AI could substitute for human lawyers in handling legal matters, when the company had not actually tested its AI’s ability to provide legal assistance and did not employ any attorneys.

In September 2024, the FTC announced Operation AI Comply, a law-enforcement initiative targeting deceptive claims about AI products.<sup>41</sup> The agency issued warning letters to multiple companies and filed complaints against firms that falsely claimed their products used advanced AI technology, or that exaggerated their AI systems’ capabilities.<sup>42</sup>

---

<sup>37</sup> 49 USC § 41713.

<sup>38</sup> Accountability Comments 15.

<sup>39</sup> 15 U.S.C. § 45.

<sup>40</sup> Press Release, *FTC Finalizes Order with DoNotPay That Prohibits Deceptive “AI Lawyer” Claims, Imposes Monetary Relief, and Requires Notice to Past Subscribers*, FED. TRADE COMM’N (Feb. 11, 2025), <https://www.ftc.gov/news-events/news/press-releases/2025/02/ftc-finalizes-order-donotpay-prohibits-deceptive-ai-lawyer-claims-imposes-monetary-relief-requires>.

<sup>41</sup> Michael D. Meuti & Andrew J. Jarzyna, *One Year In, FTC’s “Operation AI Comply” Continues Under New Administration, Signaling Enduring Enforcement Focus*, BENESCH (Oct. 21, 2025), <https://www.beneschlaw.com/resources/one-year-in-ftcs-operation-ai-comply-continues-under-new-administration-signaling-enduring-enforcement-focus.html>.

<sup>42</sup> Press Release, *FTC Announces Crackdown on Deceptive AI Claims and Schemes*, FED. TRADE COMM’N (Sep. 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>.

A technology-neutral framework like Section 5 is economically efficient because it focuses on outcomes (e.g., whether the consumer was deceived or injured) rather than processes (e.g., whether the company used an algorithm).

## **B. Civil-Rights Law and Algorithmic Discrimination**

Civil-rights law provides further examples of the efficacy of technology-neutral regulations. The Civil Rights Act of 1964 (particularly Title VII, governing employment discrimination) and the Fair Housing Act both address discrimination without regard to decisionmaking methods. These statutes prohibit discriminatory outcomes, not discriminatory technologies.

The EEOC issued technical assistance in May 2023 clarifying that Title VII and the Americans with Disabilities Act both extend to employers' use of algorithmic hiring tools.<sup>43</sup> In September 2023, the EEOC announced a settlement with iTutorGroup for \$365,000, resolving allegations that the company programmed its recruiting software to automatically reject applicants over age 55 and certain ages under 55.<sup>44</sup>

The U.S. Justice Department (DOJ) has pursued similar cases under the Fair Housing Act. In 2023, the department filed a statement of interest in *Louis et al. v. SafeRent et al.*, supporting the plaintiffs' claim that an algorithmic tenant-screening system violated the Fair Housing Act by invoking factors such as credit history and nontenancy related debts, which disproportionately affected Black and Hispanic renters.<sup>45</sup>

These enforcement actions demonstrate that technology-neutral civil-rights laws already reach AI-enabled discrimination. There is no gap in legal authority that would require new AI-specific anti-discrimination statutes. Rather than new laws, what's needed is clearer guidance on compliance expectations—specifically, what testing and documentation practices will be deemed adequate to demonstrate nondiscriminatory outcomes.

---

<sup>43</sup> *Select Issues: Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964*, EQUAL EMP'T. OPP. COMM'N (May 18, 2023), <https://www.eeoc.gov/select-issues-assessing-adverse-impact-software-algorithms-and-artificial-intelligence-used>, available at <https://web.archive.org/web/20230519192227/https://www.eeoc.gov/select-issues-assessing-adverse-impact-software-algorithms-and-artificial-intelligence-used#expand>.

<sup>44</sup> Press Release, *iTutorGroup to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit*, EQUAL EMP'T. OPP. COMM'N (Sep. 11, 2023), <https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit>.

<sup>45</sup> Press Release, *U.S. Attorney's Office Files Statement of Interest in Fair Housing Act Case Alleging Unlawful Algorithm-Based Tenant Screening Practices*, DEP'T JUST. (Jan. 9, 2023), <https://www.justice.gov/usao-ma/pr/us-attorneys-office-files-statement-interest-fair-housing-act-case-alleging-unlawful>.

## VI. Direct Responses to the RFI's Specific Questions

### A. Question (i): What AI Activities Are Currently Inhibited by Federal Statutes, Regulations, or Policies?

Section III identifies several AI use cases that are currently inhibited by federal statutes, regulations, or policies:

- The FAA's pilot-certification regime, which assumes human pilots;
- The NHTSA's vehicle-safety standards, which assume human drivers;
- The FDA's premarket approval process for physical devices, which delay deployment of continuously learning diagnostic algorithms; and
- The CFPB's guidance that Regulation B applies to creditors using complex algorithms.

In addition, there are several other AI applications that are currently inhibited by federal statutes, regulations, or policies.

#### *I. Medicare reimbursement rules regarding telehealth*

Medicare reimbursement rules regarding telehealth, particularly 42 CFR § 410.78, are structurally misaligned with the technological capabilities of modern AI in health care. The regulation was designed for an era of episodic, human-to-human telemedicine, creating a framework that substitutes a virtual visit for an in-person one. Its core tenets, such as the requirement for a "two-way, real-time interactive communication" system and geographic restrictions on the patient's location, are based on a reactive-care model, in which a practitioner responds to a patient's needs during a scheduled encounter. This approach is incompatible with AI-enabled systems, which operate on a paradigm of continuous, asynchronous data analysis to provide proactive, preventative care.

AI-enabled monitoring shifts the clinical model from reactive to proactive by using algorithms to analyze vast streams of physiological data from wearables and other sensors. This can enable the detection of preclinical signals that precede acute health events. Clinical evidence demonstrates the effectiveness of this approach. AI algorithms have been demonstrated to identify atrial fibrillation with the accuracy of trained cardiologists, enabling earlier intervention to prevent strokes.<sup>46</sup> In critical care, an AI system developed at Johns Hopkins University was found to detect sepsis nearly six hours earlier than traditional methods, reducing patient mortality by 20%.<sup>47</sup> Similarly, AI-powered continuous glucose monitors help prevent acute events like diabetic ketoacidosis in

---

<sup>46</sup> See, e.g., *Delivering Advanced AI to Cardiac Monitoring*, IRHYTHM (retrieved Oct. 22, 2025), <https://www.irhythmtech.com/us/en/solutions-services/fda-cleared-ai>.

<sup>47</sup> Laura Cech, *Sepsis-Detection AI Has the Potential to Prevent Thousands of Deaths*, JOHNS HOPKINS U. (Jul. 21, 2022), <https://hub.jhu.edu/2022/07/21/artificial-intelligence-sepsis-detection>.

diabetes patients, reducing costly hospital admissions.<sup>48</sup> These technologies provide new capacities for constant vigilance, preventing the need for acute care, rather than simply replacing an office visit.

The specific provisions within the regulation and its associated billing codes create direct barriers to the adoption of these technologies. The regulation's definition of an "interactive telecommunications system" has a bias toward synchronous, real-time video communication,<sup>49</sup> leaving no clear reimbursement pathway for the value generated by an AI's asynchronous analysis. Furthermore, reimbursement is tied to the time of specific human practitioners, with no conceptual space to pay for an algorithmic service.

This regulatory misalignment generates significant economic friction, primarily through its interaction with Medicare's fee-for-service (FFS) payment system. The FFS model rewards the volume of services delivered, creating a financial disincentive for providers to adopt preventative technologies that reduce hospitalizations and procedures. A provider that invests in an AI platform that successfully keeps patients out of the hospital is penalized with lost revenue from those avoided admissions. This arrangement makes the adoption of clinically superior technology economically irrational for the provider, even as it saves the payer money and improves patient outcomes. The current rules fail to provide a sufficient counterincentive to overcome this fundamental conflict.

A durable solution would require a multi-pronged policy approach. First, 42 CFR § 410.78 must be modernized to create a distinct benefit category for asynchronous, algorithm-based monitoring services, decoupling them from the outdated synchronous-communication requirements. Second, the Centers for Medicare & Medicaid Services (CMS) should develop a new reimbursement pathway, with billing codes that compensate for the clinical value of AI-generated insights, such as early risk detection, rather than for process metrics like minutes spent on a call.

## 2. *Federal employee productivity tools*

There is a disconnect between the U.S. government's official pro-innovation stance on AI and the implementation of this policy posture at the agency level.<sup>50</sup> While the White House seeks to remove bureaucratic restrictions, agencies—including the Veterans Affairs Department and Energy Department—have instituted blanket prohibitions on generative-AI productivity tools, citing security concerns.<sup>51</sup> This contradiction stems from gaps in the guidance provided by OMB Memorandum

---

<sup>48</sup> See, e.g., Chenyang Ji, Tong Jiang, Luolin Liu, Jiale Zhang, & Liangzhen You, *Continuous Glucose Monitoring Combined with Artificial Intelligence: Redefining the Pathway for Prediabetes Management*, 16 FRONT. ENDOCRINOL. 1571362 (May 26, 2025).

<sup>49</sup> 42 CFR § 410.78(a)(3) ("Interactive telecommunications system means... multimedia communications equipment that includes, at a minimum, audio and video equipment permitting two-way, real-time interactive communication between the patient and distant site physician or practitioner.")

<sup>50</sup> Frances M. Green, *New Federal Agency Policies and Protocols for Artificial Intelligence Utilization and Procurement Can Provide Useful Guidance for Private Entities*, WORKFORCE BULLETIN (Apr. 25, 2025), <https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities>.

<sup>51</sup> Rebecca Heilweil, *More Federal Agencies Join in Temporarily Blocking or Banning ChatGPT*, FEDSCOOP (Jan. 9, 2024), <https://fedscoop.com/more-federal-agencies-join-in-temporarily-blocking-or-banning-chatgpt>.

M-24-10, issued in 2024. While the memorandum provided a detailed governance framework for “safety-impacting” and “rights-impacting” AI, it offered no clear pathway to approve low-risk productivity software. In the absence of specific guidance, some agency officials have defaulted to the most conservative position of prohibition.

The opportunity cost of this inaction is substantial. Economic research from the Federal Reserve Bank of St. Louis suggests that adoption of generative-AI tools could increase aggregate productivity by 1.1%, with the average user saving 2.2 hours per week.<sup>52</sup> Studies on specific tasks common in government work show that AI tools accelerate writing by 40% while improving quality by 18%, and programmers using AI assistants complete tasks 55.8% faster.<sup>53</sup> By forgoing these tools, federal agencies are sacrificing potential gains in government efficiency and operational effectiveness.

The current patchwork of agency-level prohibitions is a strategically incoherent application of federal policy. These bans collapse the tiered, risk-based structure intended by OMB M-24-10, whose text implies that not all AI systems warrant the same intensive governance. They also run afoul of what the Brookings Institution has observed is a bipartisan consensus that regulatory efforts should focus on high-impact systems.<sup>54</sup> Furthermore, these prohibitions create a “shadow AI” problem, in which motivated employees will turn to personal devices and commercial accounts to use these tools, moving the activity outside the government’s security perimeter and creating the unmonitored data-handling risks the policies were designed to prevent.

A more effective approach requires the OMB to issue supplemental guidance that formally defines and creates a separate governance track for low-risk AI-productivity tools. Clarifying that simple productivity tools are not subject to the most stringent requirements of M-24-10 would enable agencies to allocate their finite governance resources to those use cases already identified as potentially implicating safety concerns or protected rights.

The practical implementation of this policy would involve transitioning from prohibition to a system of secure application “allowlisting.” This cybersecurity best practice, endorsed by the Cybersecurity and Infrastructure Security Agency, operates on a “default deny” principle, where only explicitly vetted and approved software is permitted to run.<sup>55</sup> The technical feasibility of this model has already been demonstrated by successful pilot projects in secure federal environments, including the

---

<sup>52</sup> Alexander Bick, Adam Blandin, & David Deming, *The Impact of Generative AI on Work Productivity*, FED. R. BANK ST. LOUIS (Feb. 27, 2025), <https://www.stlouisfed.org/on-the-economy/2025/feb/impact-generative-ai-work-productivity>.

<sup>53</sup> Shakked Noy & Whitney Zhang, *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*, 381 SCIENCE 187 (Jul. 13, 2023).

<sup>54</sup> Sorelle Friedler & Andrew D. Selbst, *5 Points of Bipartisan Agreement on How to Regulate AI*, BROOKINGS INST. (Aug. 15, 2025), <https://www.brookings.edu/articles/five-points-of-bipartisan-agreement-on-how-to-regulate-ai>.

<sup>55</sup> *What the Heck Is Application Allowlisting in CMMC?*, TOTEM.TECH (Apr. 25, 2025), <https://www.totem.tech/application-allowlisting-cmmc>.

Department of Homeland Security’s (DHS) “DHSSChat” tool,<sup>56</sup> the FDA’s “Elsa” platform,<sup>57</sup> and the U.S. Army’s “#CalibrateAI” program.<sup>58</sup> These examples illustrate that the primary barrier to adoption is a lack of policy clarity, not a lack of secure technical solutions.

By combining a default-deny architecture with a structured allowlist of pre-approved tools, agencies could safely expand access to low-risk AI applications—such as summarization, document drafting, or data-sorting assistants—without triggering the compliance overhead intended for high-impact systems. This would unlock measurable efficiency gains while maintaining a strong security posture, aligning governance with practical risk management, rather than categorical prohibition.

### **B. Question (ii): What Specific Federal Statutes, Regulations, or Policies Present Barriers?**

Sections III and VI.A identify several specific federal statutes, regulations, or policies that present barriers:

- **49 U.S.C. § 44703** requires individuals to hold airman certificates meeting specified knowledge, physical fitness, and character standards. The statute’s definition of “airman” as an individual prevents certification of autonomous aircraft systems.
- **49 CFR § 571.114** requires theft-protection systems, including key-locking steering columns. The standard assumes a human-operated vehicle and is structurally incompatible with autonomous vehicles without steering wheels.
- **21 CFR § 814** establishes premarket-approval requirements for Class III medical devices based on a model of fixed device designs. The regulation does not accommodate continuously learning algorithms without extensive change-control submissions.
- **12 CFR § 1002.9** (Regulation B implementing ECOA) requires creditors to provide specific reasons for adverse credit actions. The CFPB has interpreted this to require more detailed explanations than some noninterpretable models can easily generate.
- Medicare reimbursement rules regarding telehealth, particularly **42 CFR § 410.78**, are structurally misaligned with the technological capabilities of modern AI in health care.

### **C. Question (iii): Where Are Administrative Tools Available but Underutilized?**

Section III identifies several administrative tools that are available but underutilized:

---

<sup>56</sup> Michael Boyce, *DHS’s Responsible Use of Generative AI Tools*, DEP’T HOMELAND SEC. (Dec. 17, 2024), <https://www.dhs.gov/archive/news/2024/12/17/dhss-responsible-use-generative-ai-tools>.

<sup>57</sup> Press Release, *FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People*, FOOD & DRUG ADMIN. (Jun. 2, 2025), <https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people>.

<sup>58</sup> Jennifer Swanson, *#CalibrateAI/Project Athena Update*, U.S. ARMY (Nov. 19, 2024), [https://www.army.mil/article/281451/calibrateaiproject\\_athena\\_update](https://www.army.mil/article/281451/calibrateaiproject_athena_update).

- **FAA waivers under 14 CFR § 107.200:** The regulation permits waivers of Part 107 requirements for unmanned aircraft operations if the applicant demonstrates safety, but the processing times to obtain these waivers average 90 days. Greater use of categorical waivers for low-risk operations could accelerate beneficial deployments.
- **FDA expedited review pathways under 21 U.S.C. § 360e-3:** The Food, Drug, and Cosmetic Act authorizes the FDA to use expedited procedures for breakthrough devices. The agency could apply these pathways more systematically to AI-enabled diagnostic and monitoring systems.
- **NHTSA temporary exemptions under 49 U.S.C. § 30113:** The statute permits exemptions from Federal Motor Vehicle Safety Standards for up to two years (extendable to three) for vehicles with safety levels at least equal to the standard's level. NHTSA could use this authority more proactively for autonomous-vehicle testing.

In addition, the CFPB, Federal Deposit Insurance Corporation (FDIC), and Office of the Comptroller of the Currency (OCC) have authority to approve pilot programs—sometimes referred to as “regulatory sandboxes”—for innovative financial products. Creating more formalized sandbox programs with published entry criteria would reduce uncertainty.

#### **D. Question (iv): What Modifications Would Enable Lawful Deployment While Preserving Regulatory Objectives?**

Sections III and VI.A identify several modifications that would enable lawful deployment while preserving regulatory objectives:

- **Transportation:**
  - Amend 49 U.S.C. § 40102 to define “operator in command” to include both human pilots and autonomous systems and direct the FAA to establish performance-based airworthiness standards for autonomous aircraft that focus on perception, decisionmaking, communication, and fault tolerance.
  - Amend 49 U.S.C. § 30102 to define “driver” to include automated driving systems capable of performing the entire dynamic-driving task, and direct NHTSA to establish safety-performance standards for such systems that address operational design domains, object and event detection and response, fallback mechanisms, and crashworthiness.
- **Health Care:**
  - Amend 21 U.S.C. § 360c to authorize the FDA to create a provisional-approval pathway for AI-enabled medical devices, permitting initial market entry based on promising (but limited) clinical data, coupled with mandatory post-market monitoring, outcome tracking, and confirmatory studies.
  - Amend 42 CFR § 410.78 to authorize Medicare reimbursement for remote patient monitoring using AI-enabled continuous-monitoring systems when clinical evidence

demonstrates equal or superior outcomes relative to traditional scheduled consultations, with payment rates reflecting the monitoring's clinical value.

- **Financial services:** Amend 15 U.S.C. § 1691(a) to specify that, when creditors use algorithmic models, the model as a whole constitutes a “policy or practice” for disparate-impact analysis, and establish that creditors satisfy adverse-action notice requirements by providing explanations that meet specified documentation and testing standards, even if the model is not fully interpretable.
- **Federal AI use:** OMB should issue clarifying guidance distinguishing between AI systems that require special governance (those making consequential decisions about individuals or affecting safety) and productivity tools that can be used with standard information-security controls. Agencies should maintain allowlists of vetted productivity tools and usage protocols, rather than categorical prohibitions.

#### **E. Question (v): What Forms of Clarification Would Be Most Effective?**

Effective forms of clarification include interagency guidance, sector-specific standards, enforcement priority statements, and conditional safe harbors. Together, these measures would give firms clearer expectations, while allowing regulators to adapt as AI evolves.

- **Interagency guidance on technology-neutral enforcement:** OSTP should coordinate an interagency working group that includes the FTC, DOJ, EEOC, and sector-specific regulators to develop unified guidance on how existing statutes apply to AI systems. The guidance should emphasize outcomes over process and provide safe harbors for specified testing and monitoring practices.
- **Sector-specific standards for adequate documentation:** Agencies should publish model documentation templates aligned with the National Institute of Standards and Technology (NIST) AI Risk Management Framework.<sup>59</sup> For example, the CFPB could specify what testing and monitoring for algorithmic bias would be deemed adequate to satisfy fair-lending obligations. EEOC could identify acceptable statistical tests for adverse-impact analysis.
- **Enforcement priority statements:** When statutory language plausibly covers AI but the application of that language remains uncertain, agencies should issue enforcement-priority statements clarifying what conduct they will and will not pursue. For example, the FTC could specify what types of AI-generated-content disclosures it considers adequate to avoid deception claims.
- **Conditional safe harbors:** Agencies could establish safe harbors stating that firms following specified practices will not face enforcement action for specified violations. For instance, FDA could create a safe harbor for AI-related medical-device updates that follow

---

<sup>59</sup> *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NAT'L INST. STDS. & TECH. (Jan. 2023), available at <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

predetermined change-control plans, stating that such updates will not be deemed to require new premarket approval.

- **Published evaluation criteria for pilots and waivers:** Agencies should publish detailed criteria explaining what evidence applicants must provide to obtain pilot-program approvals or waivers. Criteria should specify acceptable forms of safety demonstration, monitoring protocols, and risk-mitigation measures.

#### **F. Question (vi): How Might Federal Action Address Organizational Barriers?**

To strengthen federal capacity for consistent and credible AI oversight, agencies need shared technical resources and evaluation practices, which would include the following actions:

- **Create an OSTP-led AI-evaluation support center:** OSTP should establish a center that provides agencies with technical assistance to evaluate AI pilot proposals and waiver requests. The center would maintain evaluation templates aligned with NIST frameworks and host a repository of successful pilot-project case studies.
- **Establish communities of practice:** OSTP should convene regular meetings of agency AI coordinators to share lessons learned, discuss evaluation challenges, and develop common approaches to recurring problems.
- **Develop model pilot-program terms:** OSTP should work with agencies to develop model legal terms for AI pilot programs, including entry criteria, participant obligations, data-sharing requirements, monitoring protocols, and evaluation metrics.

### **VII. Recommendations**

Based on the foregoing analysis, ICLE respectfully recommends the following actions:

#### **A. Immediate Administrative Actions**

To ensure federal oversight keeps pace with rapid advances in AI, agencies need systematic updates to outdated rules and coordinated guidance on enforcement and evaluation. This section outlines a set of institutional reforms—ranging from a federal AI-rule audit to an OSTP-led evaluation support center and innovation-impact assessments—to modernize regulatory frameworks, while preserving accountability and promoting innovation.

- **Conduct a federal AI-rule audit:** OSTP should direct agencies to inventory regulations that embed assumptions about human operation, static design, or other characteristics incompatible with AI systems. The audit should prioritize high-impact sectors and identify specific regulatory provisions for modernization.
- **Issue interagency guidance on technology-neutral enforcement:** OSTP should convene FTC, DOJ, EEOC, CFPB, and sector regulators to develop unified guidance on how existing

statutes apply to AI. The guidance should emphasize that laws prohibiting fraud, discrimination, and unfair practices apply to automated systems.

- **Publish agency playbooks for pilots and waivers:** OSTP should work with agencies to develop and publish detailed guidance on obtaining approval for AI pilot programs and waivers. Playbooks should include entry criteria, model application templates, required evidence and testing protocols, and monitoring requirements.
- **Establish an AI-evaluation support center:** OSTP should create an interagency center that provides technical assistance for evaluating AI systems, maintains repositories of successful pilots and model terms, develops training materials for agency staff, and coordinates evaluation methodologies across agencies.
- **Adopt innovation impact assessment:** OMB should revise Circular A-4 to require agencies conducting cost-benefit analysis for AI-related rules to systematically assess and document innovation impacts, including effects on market entry, dynamic competition, and international competitiveness.

## **B. Congressional Actions**

Federal legislation will be essential to provide consistent, adaptive governance for AI across sectors. This section outlines key congressional actions—targeted preemption, performance-based approval pathways, and regulatory sunset provisions—to harmonize oversight, enable innovation, and ensure accountability as AI technologies evolve.

- **Enact targeted AI preemption legislation:** Congress should pass legislation prohibiting states from regulating the design, development, training, or validation of AI models, while preserving state authority to enforce general laws against fraud, discrimination, and safety violations.
- **Authorize performance-based approval pathways:** Congress should amend relevant statutes to permit agencies to approve adaptive systems through performance-based standards, as discussed in Sections III and VI.B.
- **Require regulatory sunset and review:** Congress should enact legislation requiring that any new AI-specific regulation include either a sunset provision or mandatory periodic review to assess whether the regulation's benefits continue to exceed its costs as the technologies evolve.

## **VIII. Conclusion**

AI is a general-purpose technology whose applications and potential risks remain incompletely understood. In this circumstance of high uncertainty, regulatory forbearance—defined as maintaining technology-neutral enforcement of existing law while gathering empirical evidence—is likely to produce better welfare outcomes than premature, prescriptive mandates. Existing statutes governing fraud, product safety, discrimination, and unfair commercial practices already reach

conduct enabled by AI. The primary challenge is not a legal authority gap but regulatory mismatch and market fragmentation.

Many existing federal regulations embed assumptions about human operation or static product design, creating barriers to beneficial AI deployment. In most cases, these barriers can be addressed through administrative flexibility—such as waivers, pilot programs, conditional approvals, and updated guidance—rather than requiring new statutes.

A distinct threat arises from the proliferation of conflicting state laws. A compliance patchwork imposes costs that grow multiplicatively with jurisdictional divergence; creates barriers to entry that favor incumbents; and allows restrictive states to set *de facto* national policy through extraterritorial effects. Federal preemption of state laws that regulate the design and development of AI models, paired with technology-neutral federal enforcement, would restore constitutional and economic coherence to the national market.

Rather than creating duplicative AI-specific mandates, agencies should clarify how these technology-neutral laws apply to novel systems. Agencies should also establish safe harbors for firms that follow specified testing and monitoring practices and build internal evaluation capacity. This regulatory approach permits responsible innovation while addressing demonstrated harms and preserves U.S. leadership in a technology central to future economic and strategic competition.