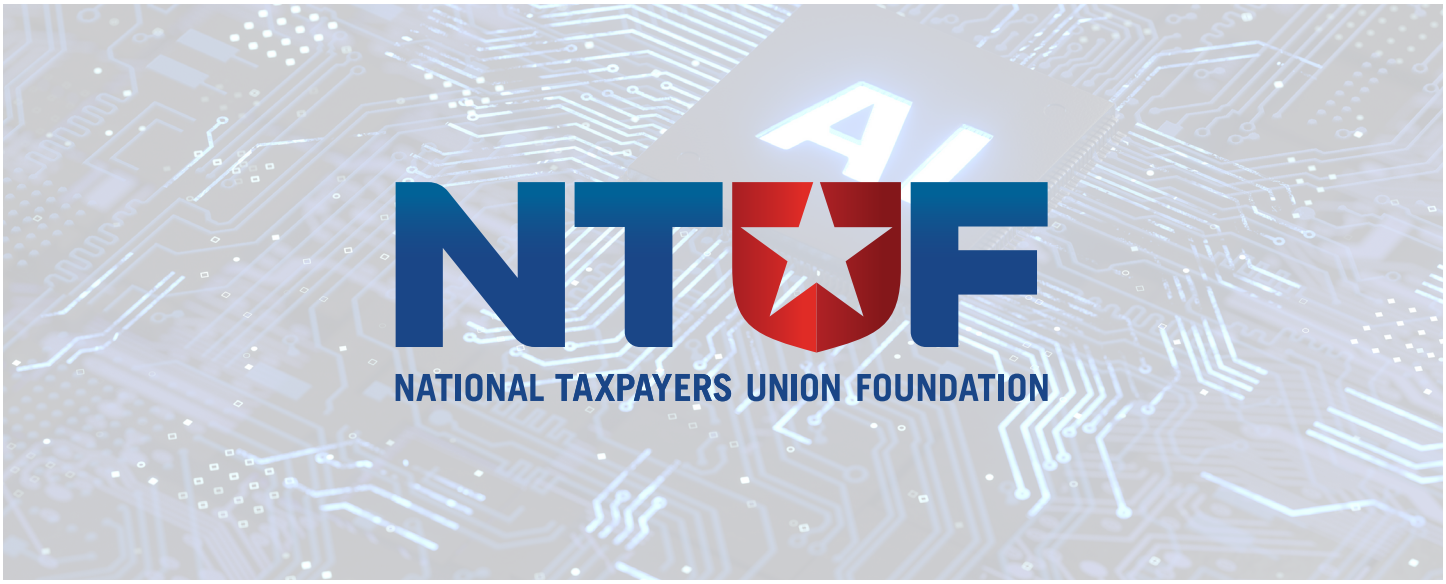

Why the United States Needs Better-Designed AI Sandboxes

RYAN NABIL
Director and Senior Fellow, Technology Policy

OCTOBER 15, 2025

Key Takeaways

- Regulatory sandboxes can be an effective tool for AI governance—but only if they are supported by the right statutory objectives and a coherent institutional design.
- The core purpose of an AI sandbox is to facilitate regulatory experimentation and learning that inform evidence-based rulemaking and reform, giving agencies a controlled environment in which to test and observe how emerging applications interact with existing or proposed rules.
- The recently introduced SANDBOX Act risks falling short of this purpose by focusing too narrowly on job creation, centralizing regulatory authority within the White House, authorizing overly broad and lengthy waivers without a clear rationale, and failing to establish mechanisms for translating sandbox lessons into systemic reform.
- A well-designed sandbox framework would strengthen institutional capacity, safeguard consumers, and ensure that the U.S. approach to AI governance evolves in response to evidence—advancing innovation while protecting the public interest.

Proposed Reforms

This brief proposes five specific reforms to improve the AI sandbox proposals:

1. Reframing the statutory purpose toward systemic regulatory reform
2. Designing multiple sector-specific AI sandboxes instead of a centralized, one-size-fits-all model
3. Limiting the scope and duration of waivers to reduce risks of regulatory privilege
4. Establishing transparent eligibility requirements and selection criteria to ensure fairness
5. Translating sandbox insights into systemic regulatory reform

Introduction

As AI applications become more advanced, regulators and lawmakers face growing pressure to craft well-calibrated rules that enable innovation while managing emerging risks. One tool borrowed from financial technology regulation—the regulatory sandbox—can be useful in that effort, but only if it is designed with the right objectives in mind.

The core purpose of an AI sandbox is not merely to promote growth but to facilitate structured experimentation and regulatory learning that inform evidence-based rulemaking and reform. More specifically, sandboxes provide a controlled environment in which regulators can test and observe how emerging AI applications interact with existing or proposed rules and determine whether those rules should be recalibrated in light of sandbox evidence. The recently introduced SANDBOX Act provides the latest opportunity to evaluate an AI sandbox policy.

What Are Regulatory Sandboxes?

Regulatory sandboxes are public-sector programs that allow startups, technology firms, and other entities to test innovative products and services under close regulatory supervision for a limited period. During the sandbox testing phase, companies may receive regulatory guidance, expedited authorization, or targeted waivers. By closely monitoring participating firms, regulators gain firsthand insights into emerging technologies and how they interact with existing or proposed rules. These insights can then serve as the basis for recalibrating regulations, eliminating outdated requirements, or developing new statutory frameworks.

Originally developed in the context of financial services, the United Kingdom's Financial Conduct Authority (FCA) launched the world's first fintech sandbox in 2016. Since then, more than 50 jurisdictions around the world have introduced similar fintech sandboxes. In recent years, policymakers and regulators have increasingly recognized the potential of regulatory sandboxes for AI governance. As a result, a growing number of jurisdictions have established or proposed AI-focused sandboxes. For example, AI sandbox proposals have been put forward in Brazil, Norway, Singapore, Switzerland, and the United Kingdom, while the European Union's Artificial Intelligence Act requires each EU Member State to create or join at least one AI regulatory sandbox. As the United States considers how best to design its own AI sandbox framework, these international experiences offer useful lessons for identifying which approaches may be most appropriate for the U.S. regulatory context.[1]

[1] For a detailed discussion of AI regulatory sandboxes in major jurisdictions and their regulatory designs, see Ryan Nabil, "Artificial Intelligence Regulatory Sandboxes, *Journal of Law, Economics, and Policy* 19, no. 2 (2024): 295–348, https://www.jlep.net/s/Nabil-Final-for-PDF.pdf.

Recent U.S. Experience with Regulatory Sandboxes

Due to overlapping federal and state regulatory authority, both federal agencies and state governments can create sandbox programs in the United States. In the context of financial services, while the Consumer Financial Protection Bureau (CFPB) launched the Compliance Assistance Sandbox, Trial Disclosure Sandbox, and No-Action Letter Policy, these initiatives did not exhibit the defining characteristics of regulatory sandboxes as commonly understood and designed in most jurisdictions. These programs were largely geared toward providing case-specific approvals and safe-harbor assurances rather than using sandbox insights to support evidence-based rulemaking and reform. Nonetheless, they remained limited in scope, and both the Compliance Assistance Sandbox and the No-Action Letter Policy were discontinued under the Biden administration in 2022. Under the Trump administration, the CFPB announced in early 2025 that it would begin accepting applications again for the Compliance Assistance Sandbox and resume issuing No-Action Letters—marking a significant shift from the approach of the previous administration.

At the state level, at least ten U.S. states created fintech sandboxes, although most of these programs struggled to attract applicants. Arizona's Fintech Sandbox and Hawaii's Digital Currency Innovation Lab were notable exceptions, attracting significantly more participants than most other state programs. However, several U.S. fintech sandbox programs had admitted no participants by late 2021, according to a Competitive Enterprise Institute (CEI) study.[2] In contrast, during the same period, the Hong Kong Monetary Authority's Fintech Supervisory Sandbox had accepted over 200 firms, while the UK FCA's Regulatory Sandbox had admitted more than 100 participants.

In the United States, the most notable regulatory sandbox has come not from fintech but from legal services. In August 2020, the Office of Legal Services Innovation of the Utah Supreme Court launched a sandbox permitting non-lawyer-owned firms and certain non-legal entities to provide services such as completing marriage, business, and immigration forms. Within a year, the program had accepted 31 firms into the sandbox—a significantly higher level of participation than in most state-level fintech programs at that time. By establishing clearer eligibility rules, allowing broader participation, and maintaining sustained regulatory engagement, Utah enabled a level of private-sector participation and innovation largely absent from other U.S. sandbox initiatives.

Since then, a number of state governments have also sought to develop AI-focused sandboxes—but their success in attracting participants and translating sandbox insights into evidence-based regulatory reform remains to be seen. In the U.S. context, the contrast between sandbox models underscores an important point: regulatory design is crucial to determining long-term effectiveness. Well-designed sandboxes can expand access to services and generate meaningful regulatory insights, whereas poorly conceived ones remain underused and may pose risks to consumers and competition. These lessons are especially relevant as policymakers consider how to design AI sandboxes that can both foster innovation and safeguard the public interest.

Key Issues with the SANDBOX Act

The central purpose of an AI regulatory sandbox is to serve as a mechanism for iterative, evidence-based rulemaking—helping regulators understand emerging applications and develop proportionate, well-calibrated rules. The SANDBOX Act, however, is not designed to generate the insights needed to support broader regulatory learning. By granting exemptions only to firms admitted into the sandbox—while leaving similarly situated firms outside still subject to the same requirements—it risks encouraging regulatory arbitrage and government-conferred privilege

[2] As of November 2021, Arizona's fintech sandbox had admitted 11 firms, while Hawaii's Digital Currency Sandbox had admitted 16, according to publicly available information and correspondence with relevant authorities. See Ryan Nabil, "How Regulatory Sandbox Programs Can Promote Technological Innovation and Consumer Welfare," *Competitive Enterprise Institute OnPoint*, no. 281 (August 17, 2022), https://cei.org/studies/how-regulatory-sandbox-programs-can-promote-technological-innovation-and-consumer-welfare/.

for selected participants.[3] In practice, this would reward firms with access while disadvantaging competitors, thereby entrenching competitive inequality and undermining the very purpose of a sandbox: to inform system-wide regulatory improvement, not to confer firm-specific benefits. The key challenges in the legislation—and potential avenues for addressing them—are discussed below.

Guidelines for Reform

The first challenge with the proposed legislation is that, instead of emphasizing regulatory learning and reform, the Act defines the purpose of a national sandbox primarily in terms of job creation and economic growth. While employment and growth are valuable policy objectives, sandboxes are not designed to operate as industrial policy tools. Their central function is to support iterative, evidence-based rulemaking—helping regulators strengthen institutional capacity, protect consumers, and ensure that rules evolve in response to evidence. By observing firsthand how specific AI applications interact with existing or proposed regulations, policymakers can calibrate regulatory frameworks that achieve a more effective balance between innovation and safety. Without this focus, a sandbox risks becoming an instrument of short-term industrial policy rather than a mechanism for generating the insights necessary for durable regulatory improvement.

A more effective approach would be to place regulatory learning and evidence-based rulemaking at the core of the AI sandbox's design. The specific features of the sandbox—including the choice of regulators, the criteria for admitting participants and granting regulatory waivers, and the mechanisms for reviewing and implementing lessons learned—should follow from these core objectives. Without a durable, evidence-based framework, the proposed sandbox risks being perceived as a forum for regulatory arbitrage rather than institutional learning and reform. Such a perception increases the likelihood that a future administration will unwind the program, repeating the familiar cycle in which successive administrations reverse or discontinue sandbox initiatives launched by their predecessors.

Designing Multiple Sector-Specific AI Sandboxes Instead of a Centralized, One-Size-Fits-All Approach

The SANDBOX Act proposes the creation of a centralized AI sandbox within the White House Office of Science and Technology Policy (OSTP), granting it broad authority to issue regulatory waivers. Yet the OSTP, as part of the Executive Office of the President, is not a regulatory agency and lacks both the enforcement powers and the sector-specific expertise of bodies such as the Food and Drug Administration (FDA), the Department of Transportation (DOT), the Federal Trade Commission (FTC), and the Federal Communications Commission (FCC). While regulatory design matters for fintech sandboxes, it is especially critical for AI sandboxes, because AI applications often fall under different regulators across sectors—and, in some cases, under multiple regulators simultaneously. Centralizing regulatory authority in the Executive Office of the President—and empowering its director to override sectoral regulators in the appeals process—risks sidelining the very bodies responsible for developing and enforcing AI rules within their remit. Such centralization weakens interagency coordination and limits the generation of the sector-specific insights essential for effective regulation.

A better approach would be to create multiple sandbox programs that reflect the regulatory architecture of different sectors. In financial services, where regulatory authority is fragmented across multiple federal agencies—and, in some areas, like insurance, state regulators—an AI sandbox would require a *single-sector, multi-regulator* model supported by a well-structured interagency process. In certain areas, a *single-sector, single-regulator* model may be more appropriate. For example, the FDA could oversee a sector-specific AI sandbox for medical devices, either independently or in consultation with other regulators. Sector-specific sandboxes are also essential for attracting and supervising a sufficient number of relevant projects to build expertise in highly

[3] For a detailed discussion of how best to address potential issues with regulatory privilege from sandboxes, see Brian Knight and Trace Mitchell, "The Sandbox Paradox: Balancing the Need to Facilitate Innovation with the Risk of Regulatory Privilege," *South Carolina Law Review* 72, no. 2 (2020): 445–475, https://sclawreview.org/article/the-sandbox-paradox-balancing-the-need-to-facilitate-innovation-with-the-risk-of-regulatory-privilege/.

specialized technologies and business models within each sector. Aligning sandbox design with sectoral regulatory structures is therefore more likely to avoid the pitfalls of a one-size-fits-all approach and to generate valuable insights as regulators develop and calibrate rules tailored to specific sectors.

Limiting the Scope and Duration of Regulatory Waivers to Reduce Risks of Regulatory Privilege

The SANDBOX Act would grant waivers that are both overly broad and excessively long, often without a clear rationale—potentially providing exemptions from almost any regulatory provision and extending them for up to twelve years. Such an approach undermines the primary purpose of a sandbox, which is to provide narrow, time-limited relief—granted under transparent criteria—to test rules and generate evidence for future rulemaking. Rather than encouraging regulators to translate lessons into broader reform, extended waivers risk entrenching advantages for a limited number of well-connected firms.

A more effective approach would be to ensure that waivers are narrowly tailored, temporary, and directly linked to identified regulatory shortcomings. More specifically, relief should be granted only when a particular rule materially impedes testing or deployment, and the insights gained should inform reforms that apply to all similarly situated firms. Otherwise, sandboxes risk reinforcing regulatory privilege for incumbents rather than serving as instruments of institutional learning and reform.

Sandbox testing duration should be limited to the period necessary for firms to bring products into compliance and for regulators to gather data—long enough to be meaningful, but not so long that it creates lasting advantages for participating firms. Care should be taken to ensure that the testing period extends no longer than required to meet the sandbox's legitimate policy objectives. As Hilary Allen of the Washington College of Law notes, most existing fintech sandboxes limit testing to between six months and two years. For U.S. AI sandboxes, a one- to two-year testing period—with limited discretion for regulators to adjust the duration by several months based on the nature of the product—would allow sufficient oversight and evidence gathering without entrenching regulatory privilege for selected firms.

Establishing Transparent Eligibility Requirements and Selection Criteria to Ensure Fairness and Regulatory Reform

As currently drafted, the SANDBOX Act does not establish clear eligibility requirements or selection criteria for participation. By contrast, most state-level financial technology sandbox statutes include specific product- and firm-level entry requirements and selection criteria. Without transparent, evidence-based standards, admission decisions risk becoming ad hoc or biased, favoring politically connected firms over those with genuinely innovative products more likely to generate meaningful regulatory insights. This would undermine the sandbox's purpose as a tool for institutional learning and reform and amplify the risk of regulatory privilege by conferring special advantages on a small set of firms without a clear public rationale.

A better approach would entail establishing transparent application and selection processes, with publicly available eligibility requirements and selection criteria. Eligibility rules should be broad enough to allow participation by firms from different sectors, including foreign companies seeking to test innovative products and enter the U.S. market, but final selection should prioritize projects most likely to generate useful data for regulators and identify areas where existing rules are unclear, unnecessarily burdensome, or outdated. Transparent eligibility and selection processes would help ensure fairness, prevent favoritism, and direct limited regulatory resources toward projects with the greatest potential to inform systemic reform.

Translating Sandbox Insights into Systemic Regulatory Reform

The Act does not currently provide any meaningful mechanism for converting sandbox findings into broader regulatory reform. The proposed annual reports to Congress and joint resolutions of approval are too infrequent and politically cumbersome to enable timely adjustments. The real value of a sandbox lies in generating insights that feed directly into agency rulemaking and guidance, allowing regulators to respond swiftly to emerging evidence. Without this feedback loop, exemptions granted within the sandbox will remain isolated—benefiting only participating firms while doing little to improve the regulatory framework as a whole.

To avoid this issue, sandbox insights must be systematically reviewed and applied to inform broader regulatory reform. Policymakers should conduct periodic evaluations of sandbox data and outcomes to ensure that evidence is used to revise existing rules or propose new statutory measures where appropriate. Formal review mechanisms—such as mandated evaluation cycles, periodic interagency reviews, and public reports similar to those published by the UK FCA's Regulatory Sandbox or the Monetary Authority of Singapore's FinTech Regulatory Sandbox—are essential to ensure that lessons are incorporated into broader regulatory frameworks. Unless sandbox insights are applied to the market as a whole, they risk falling short of their purpose and entrenching regulatory privilege for a small set of firms.

Conclusion

AI sandboxes are not substitutes for carefully designed AI regulations, nor should they be conceived as vehicles for short-term industrial policy. Their real value lies in helping regulators understand emerging technologies, identify where existing rules fall short, and generate insights for evidence-based rulemaking and reform. The SANDBOX Act risks missing this purpose by focusing narrowly on job creation, centralizing authority within the White House, granting overly broad and lengthy waivers without clear justification, and failing to establish meaningful mechanisms for translating sandbox insights into system-wide regulatory improvement.

Other provisions—such as the OSTP's authority to override decisions by other agencies, or the bill's reliance on consumer disclosures as safeguards—further underscore the need for a more careful regulatory approach. Without a more effective sandbox framework that commands bipartisan support, the initiative risks being discontinued by a future administration, undermining efforts to develop a durable, evidence-based, and nonpartisan foundation for U.S. AI governance.

A better-designed sandbox framework would reflect the regulatory architecture of relevant sectors, establish transparent eligibility requirements and selection criteria, impose strict limits on regulatory waivers, and ensure that sandbox insights are incorporated into broader rulemaking. By placing regulatory learning and reform at the center, U.S. policymakers can make sandboxes a key instrument of iterative, evidence-based AI governance—enabling the United States to regulate AI applications more effectively while fostering innovation and protecting consumers.

2025 National Taxpayers Union Foundation
122 C Street NW, Suite 700,
Washington, DC 20001
ntuf@ntu.org