



## Federal Preemption and AI Regulation: A Law and Economics Case for Strategic Forbearance

by Kristian Stout

The House Energy and Commerce Committee's recent [proposal](#) for a federal moratorium on state-level artificial intelligence regulations represents more than a mere jurisdictional power play—it embodies a fundamentally sound approach to technology governance. As AI continues its rapid evolution across virtually every sector of the economy, the question is not whether regulation will emerge, but whether it will emerge in a form that maximizes social welfare while preserving America's competitive advantage in this transformative technology.

From a law and economics perspective, the current trajectory toward fragmented state-level AI regulation threatens to impose substantial costs while failing to achieve the consumer protection goals that motivate such regulation. The proposed moratorium offers a strategic pause—what we might call “regulatory forbearance”—that can prevent these inefficiencies while allowing for the development of a more coherent, evidence-based national framework.

### Why Premature Regulation is Economically Hazardous

One of the most compelling arguments for regulatory restraint lies in the fundamental definitional ambiguity surrounding artificial intelligence itself. AI is not a monolithic technology but rather a diverse array of computational approaches, ranging from narrow rule-based systems to sophisticated large language models and everything in between. This definitional fluidity creates a problem of regulatory targeting.

When regulators attempt to craft rules for poorly defined categories, they inevitably face a tradeoff between [over-inclusion and under-inclusion](#). Rules that are too broad capture beneficial applications alongside potentially harmful ones, creating unnecessary compliance costs and deterring innovation in areas where AI can provide clear social benefits. Conversely, rules that are too narrow may miss emerging applications entirely, leaving regulatory gaps that undermine the intended protective purpose.

The economic cost of this regulatory mismatch is amplified by AI's characteristic as a general-purpose technology with applications across numerous industries and use cases. Unlike more narrowly focused technologies, AI's pervasive potential means that regulatory errors, whether in the form of over-regulation or under-regulation, can cascade across the entire economy. The optimal regulatory response to such uncertainty is often to wait for additional information rather than to act precipitously based on incomplete knowledge.

---

**Kristian Stout** is Director of Innovation Policy at International Center for Law & Economics (ICLE). Before practicing law, Kristian worked as a technology entrepreneur and a lecturer in the Computer Science Department at Rutgers University.

## Quantifying the Costs of Regulatory Fragmentation

The economic case for federal preemption becomes even more compelling when we examine the specific costs imposed by regulatory fragmentation. Recent data indicates that over [1,000 AI-related bills](#) have been introduced across state legislatures, creating what can only be described as a regulatory cacophony. Each proposed or enacted law represents not merely a potential compliance obligation, but also a source of legal uncertainty that can chill investment and redirect resources from productive innovation to defensive legal planning.

The compliance burden imposed by disparate state regulations exhibits characteristics that are particularly problematic for market efficiency. First, these costs are largely fixed rather than variable, meaning they do not scale proportionally with firm size or output. A startup developing an AI application faces essentially the same legal analysis and compliance infrastructure requirements as a large corporation, creating significant barriers to entry that can lead to market distortion.

Second, the costs are multiplicative rather than additive when firms operate across multiple jurisdictions. A company seeking to deploy AI services nationally cannot simply choose the most favorable regulatory regime; it must comply with the most restrictive requirements across all relevant jurisdictions at the same time. This “race to the bottom” creates a regulatory regime composed of the most stringent requirements, even when those requirements are neither harmonized nor logically related to each other.

Moreover, AI technologies often exhibit strong network effects, where the value of a system increases with the number of users (who can provide feedback into the system as they use it), the diffusion of API-enabled applications, and the breadth of data available for training and improvement. Fragmented state regulation can impede these effects by creating artificial barriers to data sharing, user acquisition, and system interoperability.

Consider, for example, how different state privacy requirements for AI training data could fragment the datasets available to developers, potentially reducing the quality and effectiveness of AI systems. Despite best efforts, data bias will inevitably result from this, meaning that the most “typical” types of data will be overwhelmingly represented in a data set while minority representations will be even further diminished.

Or consider how varying algorithmic auditing requirements could make it economically infeasible to deploy consistent AI systems across state lines. We have already seen this [dynamic at play](#) across international boundaries, where many cutting-edge AI capabilities are disabled in products deployed in the European Union.

Empirical evidence from analogous regulatory scenarios underscores the significant economic burdens of fragmented compliance regimes. Our failure to enact a federal-level privacy regime is instructive here. The lack of uniformity in US privacy laws has been projected to [impose approximately](#) \$98–\$112 billion per year in additional compliance costs, potentially surpassing \$1 trillion over a decade. Notably, small businesses alone would face annual compliance expenditures of about \$20–\$23 billion, substantially diverting resources from innovation and operational growth.

In short, heavy regulatory burdens essentially function as a tax on firms that dampens innovation incentives. One [study](#) found that regulatory compliance costs effectively act like a 2.5% tax on profits, leading to about a 5.4% reduction in aggregate innovation output in the economy. In other words, when firms must spend more effort and money on meeting fragmented or complex regulations, their capacity to invest in new products and technologies falls, slowing overall innovation rates. Fragmented rules amplify this effect by multiplying the compliance checkpoints and legal

reviews needed for each new initiative.

These fixed compliance costs disproportionately burden smaller firms and new market entrants, significantly raising barriers to entry and distorting market dynamics. If AI regulation were to follow a similarly fragmented path, startups and innovators would likely face analogous barriers, reducing market competition, stifling innovation, and limiting the potential economic and societal benefits of artificial intelligence. Moreover, these regulatory barriers don't just increase costs—they can fundamentally limit the scale at which beneficial AI applications can operate, reducing their social value.

### **Historical Precedent: Lessons from the Commercial Space Launch Amendments Act**

The economic logic underlying the proposed AI moratorium finds compelling precedent in the [Commercial Space Launch Amendments Act \(SLAA\) of 2004](#). Faced with an emerging commercial human spaceflight industry, Congress wisely implemented a regulatory learning period that restricted the Federal Aviation Administration from promulgating new safety regulations for crew and participant protection during spaceflight operations.

The SLAA moratorium was explicitly designed to prevent regulatory uncertainty from stifling a nascent industry with significant economic potential. The legislation recognized that premature regulation based on limited operational data could impose costs that exceeded benefits, particularly in an industry characterized by high upfront investments and uncertain technological pathways.

The economic success of this approach is evident in the substantial growth of the commercial space industry during the moratorium period. Companies like SpaceX, Blue Origin, and Virgin Galactic were able to develop and test their technologies without the burden of premature safety mandates that might have been based on an incomplete understanding of the relevant risks and appropriate mitigation strategies.

The parallels between commercial spaceflight in 2004 and AI development today are striking. Both involve rapidly evolving technologies with transformative economic potential, face regulatory uncertainty that could deter investment and innovation, and would benefit from a period of real-world deployment and data collection before comprehensive regulatory frameworks are established.

Like the SLAA, the proposed AI moratorium would not create a regulatory vacuum. Existing laws governing fraud, discrimination, consumer protection, and product liability would continue to apply to AI systems. The moratorium would specifically target the proliferation of AI-specific regulations that single out these technologies for special treatment, often without clear evidence that such treatment is necessary or economically justified.

### **The Technology Neutrality Principle: Efficient Regulation Through Existing Legal Frameworks**

A cornerstone of economically efficient regulation is the principle of [technology neutrality](#)—the idea that laws should [focus on outcomes](#) rather than the specific technologies used to produce those outcomes. This approach promotes allocative efficiency by avoiding government “picking of winners and losers” among competing technological approaches.

The argument for technology neutrality is particularly strong in the AI context because existing legal frameworks already provide substantial protection against the types of harms that AI systems might cause. Tort law principles of negligence and product liability can address AI systems that cause physical or economic harm. Consumer protection statutes can remedy misrepresentations about AI capabilities or unfair practices in AI deployment. Anti-discrimination laws can address biased outcomes, regardless of whether they stem from AI systems or other decision-making processes.

This existing legal toolkit has several economic advantages over AI-specific regulation. First, it has been refined through decades of judicial interpretation and enforcement, reducing the likelihood of unintended consequences or regulatory gaps. Second, it applies consistent standards across different technologies, promoting competition and preventing regulatory arbitrage. Third, it relies primarily on ex-post liability rather than ex-ante regulation, allowing market forces to drive innovation while providing remedies for actual harms.

Critics of the moratorium argue that existing laws are [insufficient](#) to address AI-specific risks. However, this argument often conflates the theoretical possibility of novel harms with the practical necessity for new regulatory tools. The relevant question is not whether AI might create new types of problems, but whether existing legal frameworks are inadequate to address those problems in a cost-effective manner.

To date, the evidence suggests that many purported AI-specific harms are actually manifestations of familiar legal problems in new technological contexts. For example, fraud—a significant concern when it comes to AI-assistance malfeasance—is nothing new. Enforcement of existing laws may require more resources to address this problem at scale. But that is a different question from whether law enforcement needs a new law in order to pursue fraudsters.

Similarly, algorithmic bias in hiring decisions is fundamentally a discrimination problem that existing civil rights laws are designed to address. The fact that discrimination occurs through automated systems rather than human decision-making does not necessarily require new legal approaches—it may simply require existing laws to be enforced more effectively.

Indeed, many of the novel problems speculated with respect to AI are on the order of [existential risks](#). Leaving aside the probability of any of these existential risks for the moment, these are, by their nature, something much more appropriate for Congress to consider as part of a holistic view of the economy, national security, and foreign policy.

### **The Commerce Clause and Interstate AI Markets**

The constitutional basis for federal preemption of state AI laws rests primarily on the Commerce Clause, which grants Congress broad authority to regulate interstate commerce. The inherently interstate nature of AI development and deployment provides a strong foundation for federal regulatory authority.

AI systems typically involve data flows, computational resources, and service delivery that cross state boundaries. Large language models are trained on datasets that aggregate information from across the country and around the world. AI services are deployed through cloud computing platforms that may process data in multiple states. Even locally deployed AI systems often rely on models or updates that originate elsewhere.

These characteristics mean that state-level AI regulation often has extraterritorial effects, regulating conduct that occurs in other states or imposing compliance costs on interstate commerce. From a constitutional law perspective, such regulations are vulnerable to dormant Commerce Clause challenges. From an economic perspective, they represent precisely the type of regulatory balkanization that the Commerce Clause was designed to prevent.

Further, the economic argument for preemption in the AI context is strengthened by the fact that the technology does not fit neatly into traditional categories of local concern. Unlike restaurants or retail establishments, AI systems often cannot be confined to a single state's jurisdiction. The network effects and scale economies that characterize AI development mean that the benefits of uniform national regulation are particularly pronounced.

## Consumer Protection and Democratic Values

Critics of the proposed moratorium raise legitimate concerns about consumer protection and the value of state-level policy experimentation. These arguments deserve serious consideration from both legal and economic perspectives.

The primary critique is that the moratorium would leave consumers vulnerable to AI-driven harm by preventing states from enacting protective regulations. This argument has both descriptive and normative components that merit analysis.

Descriptively, the argument assumes that state AI regulations would provide meaningful consumer protection. However, this assumption is questionable given the definitional challenges and technical complexities involved in AI regulation. Poorly designed state laws might create a false sense of security while imposing real economic costs, resulting in a net reduction in consumer welfare.

Normatively, the argument assumes that consumer protection requires AI-specific regulation rather than enforcement of existing laws. As discussed above, this assumption is questionable given the breadth of existing legal protection and the potential for AI-specific rules to become quickly outdated or counterproductive.

A second critique invokes the traditional role of states as “laboratories of democracy,” arguing that federal preemption would prevent valuable policy experimentation. This argument has significant force in many regulatory contexts, but its applicability to AI is limited by the technology’s inherently cross-border nature.

The economic benefits of regulatory experimentation decline when the costs of fragmentation are high and when successful experiments cannot be easily contained within state boundaries. In the AI context, both conditions are present. The high costs of compliance with multiple regulatory regimes reduce the net benefits of experimentation, while the interstate nature of AI deployment means that successful state-level innovations are likely to have national implications that warrant federal consideration.

## Welfare Effects of Alternative Regulatory Approaches

A complete analysis requires consideration of the welfare effects of different regulatory approaches. While empirical data on AI regulation remains limited, we can nonetheless theoretically consider the likely effects of fragmented state regulation versus federal preemption with regulatory forbearance.

Fragmented state regulation imposes several types of economic costs. Direct compliance costs include legal fees, technical modifications to systems, and administrative overhead. These costs are largely fixed and therefore represent a particularly heavy burden on smaller firms and new entrants.

Indirect costs include the opportunity cost of diverted resources and the dynamic effects on innovation incentives. When firms must spend resources navigating regulatory complexity rather than developing new products or services, the entire economy suffers from reduced innovation and productivity growth.

Perhaps most significantly, fragmented regulation can create barriers to entry that distort market structures. If only large firms can afford the compliance costs associated with operating in multiple jurisdictions, the result may be a less dynamic AI industry—exactly the opposite of what consumer protection and competition regulation is intended to achieve. Ironically, one of the

concerns raised at the recent House E&C hearing was that “Big Tech” is dominating AI. From our current perspective, it’s difficult to know (or perhaps ever fully know) what the ideal structure of the “AI Market” is or will be (assuming such a market can even be delineated). It is entirely possible that only a handful of large players might be able to provide enough capital to support that market. One thing is sure, however, and that is that excessive compliance costs will guarantee that the ultimate number of firms able to serve the “AI Market” will be fewer than what would otherwise be the case.

The proposed moratorium can be understood as preserving the “option value” of developing better information before implementing comprehensive AI regulation. This option value is particularly high in the AI context because the technology is evolving rapidly and the optimal regulatory approach is highly uncertain.

When the costs of regulatory errors are high and irreversible, there is value in waiting for additional information even if some beneficial regulation is delayed. The rapidly evolving nature of AI technology means that regulations implemented today may be obsolete within a few years, while the costs of premature regulation—in terms of stifled innovation and reduced competitiveness—may persist much longer.

## Conclusion

The moratorium should not be viewed as an excuse for federal inaction, but rather as an opportunity for thoughtful federal engagement with AI policy issues. During the moratorium period, federal agencies should focus on understanding AI technologies, gathering data on their societal impacts, and developing the technical expertise necessary for effective oversight.

This approach mirrors the successful model of the SLAA, which combined regulatory forbearance with active federal engagement in developing voluntary industry standards and building technical expertise within regulatory agencies.

By preventing the emergence of a costly and fragmented regulatory patchwork, the moratorium can preserve the conditions necessary for continued American leadership in AI development while allowing time for more informed and effective regulatory approaches to emerge.

The economic stakes are significant. AI represents one of the most important technological developments of our time, with the potential to drive productivity growth and improve living standards across numerous sectors. Premature or poorly designed regulation could squander this potential, while thoughtful regulatory forbearance could help ensure that the benefits of AI development are widely shared.

As Congress considers this proposal, the choice is not between regulation and deregulation, but between fragmented, premature state-level regulation and a more coherent, evidence-based approach that preserves space for innovation while maintaining essential consumer protections. The latter approach offers the best path forward for maximizing the social benefits of artificial intelligence while managing its risks responsibly.

The proposed moratorium thus represents not regulatory abdication, but regulatory wisdom—a recognition that sometimes the most productive government action is the creation of space for private innovation and learning. In the rapidly evolving field of artificial intelligence, such strategic forbearance may be the key to ensuring that American leadership in this transformative technology continues to benefit both innovation and the broader public interest.