

Independent Oversight Marketplace for AI: The Solution to the Preemption vs Patchwork Debate

Over the weekend, the AI policy debate was thrown into uncharted waters when Republicans' revised 10-year federal freeze on state AI regulation cleared the procedural "Byrd Bath" hurdle, allowing it to remain in the budget reconciliation bill. Substantial confusion about the new provision, now called a "Temporary Pause," as well as opposition from both sides of the political aisle, persists.

But regardless of its fate, the provision raises a central question that will endure beyond this legislative debate: Inseparable from the challenge of *how* to regulate AI is the question of *who* should be regulating it.

AI poses a unique governance challenge. This technology is both inescapably global in its reach – hundreds of millions of people use some form of AI today – and profoundly local in its impacts, expanding into lives, families, workplaces, and communities with differentiated effects. [Americans across party lines and geographies](#) are calling for reasonable guardrails for AI, and legislators and industry both understand the need for a smart governance framework that advances America's tech leadership in the world. But AI cuts confusingly across the jurisdictions of local, state, and federal governments, making it almost singularly difficult to know who should govern AI, and how.

In many policy circles, the prevailing impulse has been to have Congress take the first step in AI governance. But no matter how well-intentioned any government's efforts might be, the reality is that direct government regulation simply cannot keep up with the pace of innovation. Worse yet, it might suffocate it. The EU AI Act stands as a cautionary tale of what happens when we demand that the government institutions that are slowest moving, and also most removed from the real-life effects of both the technology and its governance, take charge of AI governance.

But punting to the states and localities presents just as many issues. Chief among them is the risk of creating a messy patchwork of regulations that forces AI companies to move resources



from innovation toward compliance, a switch that could prove fatal to small developers, which lack the armies of compliance lawyers that Big Tech can afford. A worst-case scenario would see the American AI ecosystem collapsing under the weight of wholly contradictory and conflicting laws that cannot possibly be reconciled. American innovation and competitiveness would grind to a devastating halt, as China – unburdened by any such regulations – leaps ahead.

As heated as this debate is, it exists within a false paradigm. There is no binary choice between federal-only and state-only, or, as it's seen in the current policy dialogue, between “preemption and patchwork.”

The time is ripe for institutional entrepreneurship in AI governance. That is why we founded Fathom, an independent nonprofit policy organization that finds, builds, and scales innovative AI governance solutions that both increase public trust in AI and unleash American innovation. We are conveners and builders who bring together individuals and organizations from across every facet of the AI universe, and then architect innovative policy solutions that look beyond Beltway battles to actually address the needs of society and innovators alike.

We believe there is a third-way solution to AI governance that will benefit both consumers and industry by making AI systems more accountable, trusted, and therefore widely used. This approach is called an “independent oversight marketplace for AI.”

This approach could operate via a range of mechanisms, but lawmakers and policy experts are increasingly supporting the enactment of Independent Verification Organizations for AI (IVOs). Under this voluntary system, a government, be it state or federal, would set goals related to AI safety and risk levels, and then authorize a marketplace of regulatory bodies called IVOs – made up of subject matter experts, independent from both industry and government – to develop technical criteria that measure if an AI company has met those outcomes.

AI companies can voluntarily pursue certification from an IVO, providing a “gold standard” seal of approval, which signals that a heightened standard of care has been met to make their products safe. A state could choose to heighten the incentive for companies to participate, potentially by authorizing the certifiers to confer tort protection, which would give industry the legal certainty needed to innovate with confidence. No matter the exact carrot, the system advantages safer, more accountable AI systems and products, allowing the most trustworthy ones to rise to the top.



The idea of an independent oversight marketplace for AI is relatively new to the AI policy conversation, but not to other formerly frontier technologies. They've helped promote innovation and achieve safety in industries like the electrical grid, train system, and much of our financial system. They have precedents outside of emerging technologies, as well: in the twentieth century, the safety science company Underwriters Laboratories performed testing and certification for consumer products ranging from lithium batteries to smoke alarms.

This approach solves several problems that neither a moratorium nor a Wild West of state governance is equipped to address.

First, it empowers technical experts. AI is poorly suited to governance by non-experts, given its complexity; governance by non-technical lawmakers risks creating heavy-handed and ineffective rules that place downward pressure on innovation. IVOs elevate technical experts in the design and certification of standards for AI development, ensuring a technical and non-arbitrary system of standards.

Second, this approach is inherently set up to move at the speed of innovation. Legislative bodies are designed to move slowly to ensure thoughtful deliberation, and their traditional modes of governance – direct regulations and licensing regimes – are too static and inflexible for a fast-evolving, all-encompassing technology. IVOs would be required to constantly evolve standards and methods as the technology develops.

Third, this marketplace relies on market forces, not government regulators, to direct attention to the areas of greatest risk. AI companies that adhere to best practices gain market advantage over competitors, turning accountability and safety into a competitive advantage.

Finally, fourth – crucially – an independent oversight marketplace invites leadership from the states and yet has the capacity to scale across the country. Much like driver's licenses and many professional licenses, each state could run its own process, while recognizing the outcomes and standards set in other states. This is not a zero-sum game; different markets could borrow from others' expertise and regulatory innovation.

If the federal government wishes to adopt this model, the state-led independent oversight marketplace provides a foundation for federal action. A national system could quickly become a federal one, with a federal authority, such as NIST, setting outcomes, rather than the states, and approving IVOs.



An independent oversight marketplace for AI won't solve every challenge posed by AI – nor does it attempt to. In fact, it explicitly allows for state and federal governments to regulate other aspects of AI, including relating to deepfakes, privacy, and algorithmic bias, as they see fit. But it does put in place a model of governance that can start – right away – to address harms, create guardrails, and increase public trust in AI, in a scalable and flexible way.