

The executive order that proved Brussels right? Why AI governance requires centralisation



Dr Barry Scannell
Partner
William Fry

I. Introduction: The forgotten distinction

American criticism of European technology regulation has followed a well-worn script for years. When the EU AI Act entered into force in August 2024, Silicon Valley executives and Washington policymakers united in warning that Brussels' command-and-control approach would strangle innovation, drive investment to more permissive jurisdictions, and ultimately cede global AI leadership to China. The transatlantic regulatory divide was framed in familiar binaries of innovation versus precaution, market-driven versus state-directed, American dynamism versus European bureaucratic sclerosis.

This framing systematically misidentified the actual regulatory divergence. The meaningful difference between U.S. and EU AI governance was never the quantum of regulation, as U.S. states were rapidly implementing comprehensive AI requirements that rivalled or exceeded Brussels' mandates in scope. The true distinction was architectural. The EU AI Act operates as a single legislative instrument uniform across 27 member states, enforced through harmonised national competent authorities coordinated by a central AI Office. American AI governance, by contrast, was radically decentralised. Individual U.S. states, which possess substantial law-making powers within the American federal system, enacted their own AI regulations. California imposed disclosure requirements, Colorado prohibited algorithmic discrimination, New York mandated employment AI audits, and Illinois restricted biometric systems. Industry warnings about regulatory patchwork applied not to Europe but to the United States itself.

The Trump administration's December 2025 Executive Order, "Ensuring a National Policy Framework for Artificial Intelligence," demolishes this architectural difference through federal pre-emption of state laws. Pre-emption is the mechanism by which higher-level law overrides lower-level law in a federal system. In the United States, when federal law conflicts with state law, federal law prevails under the Supremacy Clause of the Constitution. The Executive Order attempts to achieve this outcome by declaring state AI regulations invalid and replacing them with uniform federal standards, or more precisely, by preventing states from regulating AI at all.

What emerges is a paradox that has escaped notice in the transatlantic regulatory debate. U.S. officials spent years deriding EU centralisation whilst American states were constructing their own fragmented regime, only to pursue the same structural solution once faced with the compliance complexity that industry warned Brussels would create. The EU centralises through a comprehensive legislative framework establishing baseline protections whilst the U.S. Executive Order centralises through pre-emption designed to prevent any meaningful regulatory ceiling from forming at all.

II. The Executive Order's enforcement mechanisms

The Executive Order deploys a three-pronged strategy to achieve federal pre-emption of state AI laws. Each mechanism operates on coordinated timelines designed to create comprehensive pressure on states to abandon their regulatory efforts.

The Attorney General, who heads the Department of Justice and serves as the federal government's chief law enforcement officer, must establish an AI Litigation Task Force within 30 days of the Order's signing. This Task Force has a single mandate to challenge state AI laws in federal courts on grounds that they unconstitutionally regulate interstate commerce, are pre-empted by existing federal regulations, or violate constitutional rights including the First Amendment's protection of speech. The "dormant" Commerce Clause argument posits that state AI laws burden interstate trade. This is a judicial doctrine developed by American courts that limits states' ability to regulate commerce even when Congress, the federal legislature, has not acted. The doctrine normally requires congressional action for enforcement, making the Executive Order's reliance on it constitutionally questionable. The Task Force will coordinate with the White House AI and Crypto Advisor to identify which state laws warrant federal challenge, with Colorado's algorithmic discrimination law and California's disclosure requirements serving as explicit targets.

Within 90 days, the Commerce Secretary must publish an evaluation identifying "onerous" state AI laws that conflict with the federal policy of maintaining a "minimally burdensome" regulatory environment. This evaluation must flag laws that require AI models to "alter their truthful outputs" or compel disclosures that might violate the First Amendment. The evaluation serves as the trigger mechanism for the Order's other enforcement provisions, effectively creating a federal blacklist of state regulations deemed incompatible with the administration's approach.

The Order's financial coercion operates through the Broadband Equity, Access, and Deployment (BEAD) programme, a \$42.5 billion infrastructure initiative to expand internet access in underserved areas. States identified as having onerous AI laws will be rendered ineligible for the remaining \$21 billion in non-deployment funds, which cover digital literacy, workforce development, and related programmes rather than physical infrastructure. The Commerce Secretary must issue a policy notice within 90 days establishing this funding condition, justified by the claim that fragmented AI regulation undermines broadband deployment and network-reliant AI applications. This funding mechanism essentially tells states that they must choose between their AI consumer protection laws and federal infrastructure money. Beyond BEAD, all federal agencies must assess their discretionary grant programmes to determine whether they can condition funding on states either not enacting conflicting AI laws or agreeing not to enforce existing ones during grant performance periods.

The Federal Trade Commission and Federal Communications Commission, independent regulatory agencies that operate somewhat analogously to national regulators in EU member states, receive parallel directives to establish administrative pre-emption through existing statutory authority. The FTC Chairman must issue a policy statement within 90 days explaining how state laws requiring alterations to AI outputs are pre-empted by the federal prohibition on deceptive trade practices under the FTC Act. This novel theory reframes algorithmic bias mitigation as compelling deception because adjusting AI outputs for fairness purposes allegedly requires producing "false results". The FCC Chairman must initiate a proceeding to adopt federal AI disclosure and reporting standards that explicitly pre-empt conflicting state requirements, leveraging Section 253 of the Communications Act, which prevents states from prohibiting telecommunications services.

III. Structural mimicry with normative inversion

The Executive Order pursues the identical structural outcome as the EU AI Act through radically different constitutional means towards an inverted substantive end. Like the AI Act, it treats AI as a cross-border market requiring uniform rules, stating that a "patchwork of 50 different regulatory regimes" creates compliance complexity "particularly for start-ups". This is the same market fragmentation argument that motivated Brussels, stated with the same urgency but directed towards opposite ends.

The EU AI Act centralises to impose comprehensive regulation addressing fundamental rights, safety, and transparency through mandatory conformity assessments and enforcement mechanisms backed by fines of up to 7% of global turnover. Brussels builds a regulatory floor that member states cannot undermine.

Member states cannot weaken the protections the AI Act establishes, though in some areas they can add supplementary requirements. The Executive Order centralises to prevent state regulation that the administration characterises as “onerous” or “ideologically biased”. Washington constructs a regulatory ceiling that states cannot exceed, prohibiting states from establishing protections that go beyond minimal federal standards or, in many cases, from regulating at all.

This inversion has a specific ideological driver. The Order’s critique focuses obsessively on algorithmic bias mitigation requirements, reframing them as compelling AI systems to “produce false results” to avoid disparate impact on protected groups. Colorado’s law, for instance, requires developers of high-risk AI systems to take reasonable care to prevent algorithmic discrimination in employment, housing, healthcare, and other sensitive domains. The Order characterises this as forcing AI to generate untruthful outputs. This framing transforms civil rights compliance into consumer deception and allows the FTC to invoke its authority over unfair and deceptive practices to pre-empt state anti-discrimination law. The EU AI Act, by contrast, treats algorithmic bias mitigation as a fundamental objective requiring bias monitoring and data governance for high-risk systems under Title III.

IV. Constitutional legitimacy and the gap in democratic process

The structural similarity between the EU AI Act and the U.S. Executive Order makes their constitutional divergence more striking. The AI Act’s pre-emptive force derives from primary legislation adopted through the ordinary legislative procedure, requiring approval by both the Council representing member states and the Parliament representing citizens. This is directly analogous to how EU regulations work across policy domains. When the AI Act pre-empts German federal law or Irish regulations, it does so with authority that German and Irish representatives participated in creating through the Council, whilst European citizens participated through their directly elected Members of the European Parliament.

The U.S. Executive Order lacks comparable democratic pedigree. It is not legislation passed by Congress, the bicameral legislature composed of the Senate and House of Representatives, but rather an executive decree issued by the President alone following Congress’s explicit rejection of AI pre-emption. In July 2025, the Senate voted 99-1 to remove a state AI law moratorium from the reconciliation bill, a special legislative procedure for budget-related measures. Congressional attempts to include pre-emption language in the National Defense Authorization Act, the annual military spending legislation that typically passes with bipartisan support, similarly failed. These are not close votes reflecting legislative ambivalence but overwhelming bipartisan rejection of the very policy the Executive Order now attempts to implement unilaterally.

The Order circumvents this legislative defeat through administrative workarounds. Since an American President cannot simply declare state laws void by executive decree, the Order deploys indirect mechanisms. The DOJ Task Force will argue in federal courts that state AI laws unconstitutionally burden interstate commerce under the “dormant” Commerce Clause doctrine, but this judicial doctrine cannot be invoked as independent executive authority to pre-empt state law. Courts must find state laws invalid through case-by-case adjudication. The President cannot direct courts to reach particular conclusions. Legal scholars across the ideological spectrum assess the administration’s dormant Commerce Clause claims as “legally meritless” absent new federal legislation actually passed by Congress.

The BEAD funding threat faces its own constitutional obstacles under American federalism principles. The Supreme Court’s decisions require that funding conditions be unambiguous, related to the federal purpose of the grant programme, and not so coercive as to amount



to compulsion rather than voluntary choice. Conditioning broadband deployment funds on states abandoning unrelated AI consumer protection laws surely fails the “germane-ness” test, which requires a reasonable relationship between the condition and the programme’s purpose. The infrastructure programme was designed to expand internet access in rural and underserved areas, not to enforce AI policy preferences. The constitutional problem becomes acute because the federal government is threatening to withhold funds that Congress already appropriated for specific purposes, funds that states accepted under different conditions.

The constitutional vice is that the Executive Order achieves centralisation without the constitutional mechanism that validates centralisation in the American federal system. Congress is the body constitutionally empowered to pre-empt state law through the Supremacy Clause, which establishes that federal law is “the supreme Law of the Land” when Congress acts within its delegated powers. When Congress legislates within its Commerce Clause authority to establish uniform national standards, state laws that conflict are displaced. But Congress declined to do so repeatedly, overwhelmingly, and recently. The Order thus represents centralisation through executive assertion rather than democratic deliberation.

V. Implications

The convergence on centralisation creates significant implications for global AI governance. The U.S. can no longer credibly criticise the EU AI Act’s centralised structure because the Executive Order concedes that uniform rules are necessary. American criticism must now focus on substance, whether rules should be permissive or restrictive, rather than structure, whether authority should be centralised or decentralised. This recalibration could facilitate transatlantic regulatory cooperation, as both systems accepting centralisation allows negotiations to focus on interoperability and mutual recognition rather than fundamental architecture.

Other federal or quasi-federal systems will face analogous pressures. Canada’s provinces, Australia’s states, and emerging economies with federal structures confront the same trade-off between local experimentation and market fragmentation. The U.S.-EU convergence on centralisation, despite opposite normative goals, suggests that this is not a contingent policy choice but a structural necessity driven by AI’s technical characteristics. Governments worldwide should expect similar pressures towards centralised AI governance frameworks. The battle over AI governance shifts from whether to centralise to who centralises and through what mechanisms. The EU model of multilateral legislative process with judicial review offers democratic legitimacy but slower adaptation to technological change. The U.S. model under the Executive Order of unilateral executive action offers speed and decisiveness but constitutional fragility and democratic legitimacy deficits that will almost certainly result in protracted litigation.

VI. Conclusion

The Executive Order and the EU AI Act converge on the same structural conclusion that sub-national fragmentation is incompatible with credible AI governance. This convergence has been obscured by transatlantic rhetoric of regulatory divergence. Brussels regulates, Washington deregulates. Europe applies precaution, America embraces innovation. These framings capture real normative differences in how much protection each system believes AI requires, but they mistake substantive disagreement about regulatory intensity for structural disagreement about governmental architecture.

The regulatory race that matters is not between over-regulation and under-regulation, but between legitimate and illegitimate centralisation. The EU AI Act centralises through democratic legislative process constrained by subsidiarity and proportionality principles that require justification for EU-level action. The U.S. Executive Order centralises through executive assertion contested in courts and explicitly rejected by the legislature. One flows from constitutional mechanisms designed for such centralisation. The other circumvents them.

This moment demonstrates that AI’s technical architecture exerts structural pressure towards centralised governance that transcends normative preferences and political systems. AI systems operate across borders through cloud infrastructure, train on globally sourced data, and generate effects that spill across jurisdictional boundaries. Whether governments centralise to regulate or to prevent regulation, they centralise nonetheless. The age of AI governance through fragmented territorial experimentation is ending on both sides of the Atlantic. What remains contested is whether centralisation will be legitimate, accountable, and purposive, or instead unilateral and procedurally defective. The answer will determine not only how AI is governed, but whether the governance itself can be sustained through inevitable legal challenges and political transitions.

Dr Barry Scannell
William Fry