3 Models, 1 Problem: Why AI Regulation Is Diverging Just as Multinationals Need It to Converge

Corporate, Securities, and Governance

Technology, Privacy, and eCommerce

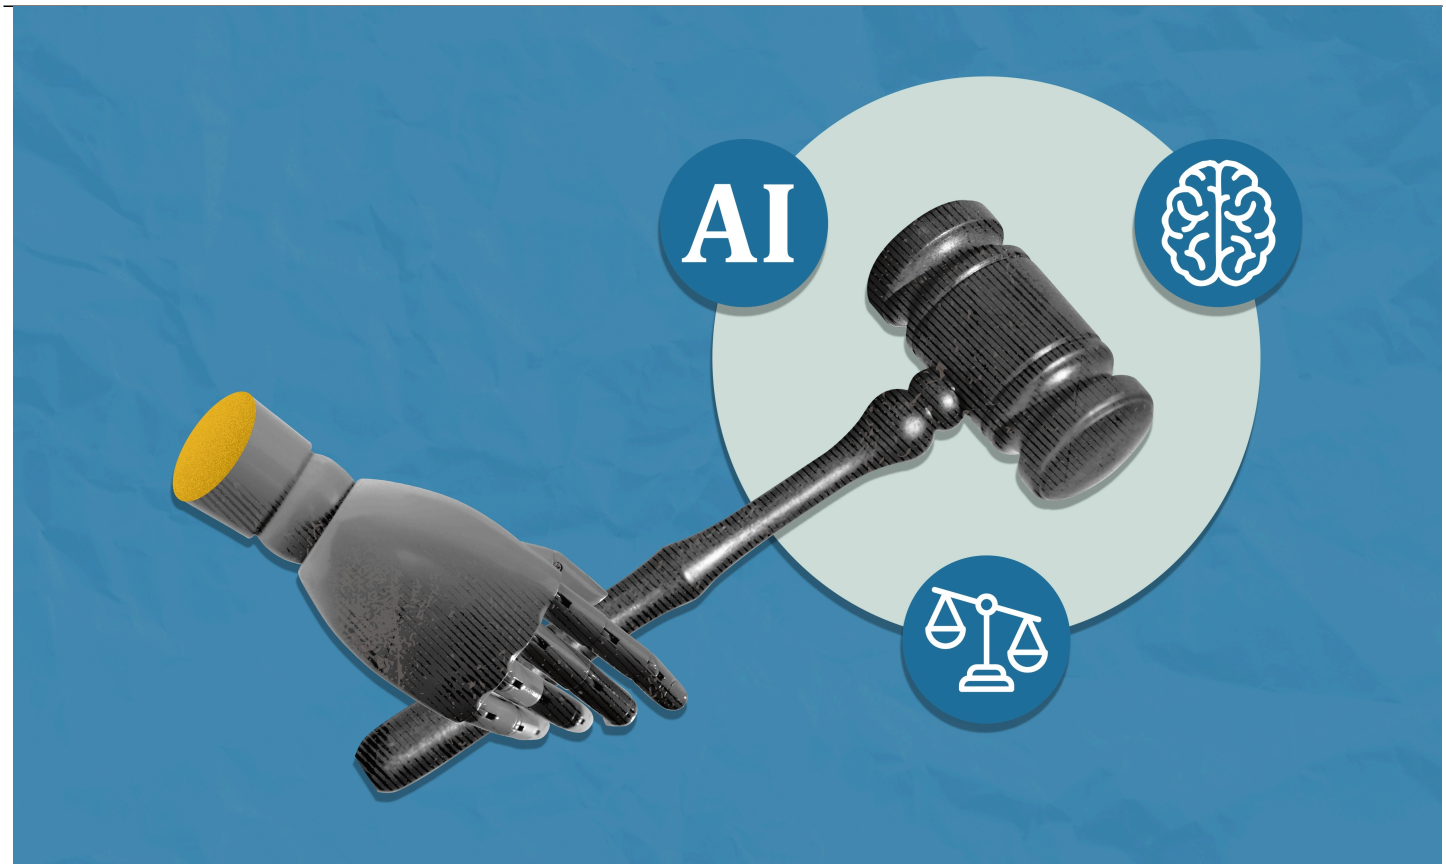

Banner artwork by / Shutterstock.com

A bank in three countries wants one answer to a simple question: is our credit scoring algorithm legal? In Brussels, the answer runs to a checklist. In Washington, there is no checklist at all, only the risk of being told after the fact that the algorithm broke a decades-old law never written with AI in mind. London provides a third answer, focused on the decision-making process and the safeguards available to individuals. Three lawyers, three countries, and three materially different pieces of advice for what is, operationally, the same piece of software.

The issue is no longer confined to technology companies. Banks, insurers, employers, retailers, and professional services firms increasingly rely on AI tools whose legal treatment now differs materially across jurisdictions.

Most compliance teams assume this is temporary, that AI regulation is simply moving at different speeds in different places and will eventually settle into roughly the same rules everywhere. That assumption is worth dropping. The EU, the US, and the UK are not on different paths to the same destination. They have built three structurally different answers to the question of what AI regulation is even for, and the gap between them is widening just as multinational clients are trying to build one global AI policy to cover all three.

What does that mean in practice for a client running the same AI system across all three jurisdictions?

In the EU, a credit scoring tool will generally fall within Annex III of the AI Act as a high-risk system. Its provider must satisfy requirements including technical documentation, conformity assessment, and registration, while the bank deploying it carries separate obligations around oversight, following the provider's instructions, and monitoring how the system performs. Compliance is largely structured and assessed in advance. Article 50 transparency obligations, requiring disclosure when users interact with AI systems and marking of AI-generated content, came into force on 2 August 2026 as scheduled. The Digital Omnibus, formally adopted in July 2026, deferred the high-risk obligations for standalone Annex III systems to December 2027, giving organizations more time to prepare. The structure of the framework, however, has not changed.

In the US, there is no comprehensive federal AI statute comparable to the EU AI Act, and as of mid-2026 no federal bill close to becoming one. Instead, older laws get stretched to cover AI conduct after the fact: the Equal Credit Opportunity Act (ECOA), which requires lenders to explain adverse credit decisions regardless of whether a human or an algorithm made them, and the broader UDAAP standard (Unfair, Deceptive or Abusive Acts and Practices), which may apply where the system produces unfair, deceptive, or abusive outcomes. There is no pre-approval process, so the legal boundary often becomes clear only through enforcement or supervisory action.

In the UK, the answer is a third thing again. The Data (Use and Access) Act 2025, whose principal data protection reforms took effect in February 2026, broadened the circumstances in which organizations may make significant decisions solely through automated processing, that is, decisions taken without meaningful human involvement. Where the rules apply, organizations must provide safeguards that let the affected person get information about the decision, make representations, challenge it, and seek genuine human intervention. The regime does not classify the system itself by risk, and there is no conformity assessment or registration requirement. Instead, it focuses on how the automated decision is taken and what protections the individual receives. The UK's own 2023 AI White Paper approach, principles enforced sector by sector rather than through one statute, remains the wider policy, and ministers have said repeatedly through 2026 that a UK AI Act is not coming any time soon.

Same algorithm, three jurisdictions, three different legal questions. For multinational groups, that may mean applying three separate governance analyses to what is operationally a single model. That has two practical consequences clients rarely see coming.

First, most in-house AI governance policies are built around the EU AI Act, because it is the most detailed and the easiest to turn into a checklist. But an EU risk assessment does not tell you whether the client's US adverse-action letters satisfy ECOA, or whether its UK process gives customers a genuine right to contest, and those are different questions the EU paperwork was never designed to answer.

Second, the same gap sits inside AI vendor contracts. A vendor agreement written with the EU AI Act in mind covers technical documentation and the risk of the system being reclassified later. It usually says nothing about who supports the client through a US adverse-action dispute, or who provides the human reviewer a UK customer is legally entitled to ask for. It may also say little about audit rights, data access, record-keeping or responsibility allocation where the different regulatory requirements collide. That is a real contractual hole, and the moment to close it is during contracting, not after a regulatory investigation or a customer dispute exposes it.

None of this is a case for waiting until the law settles down. It will not, because there is no single finish line these three systems are all racing toward. The EU regulates designated uses of the technology itself, before it is ever deployed, and has now extended that timeline through the Digital Omnibus without changing the underlying architecture. The US principally applies existing conduct-based laws, often through enforcement after deployment. The UK relies on sector regulators and safeguards built around individual automated decisions. These are not three drafts of the same regime. For in-house counsel advising multinational clients, the practical task is not to stretch one global policy across all three, it is to identify where each jurisdiction is actually asking a different question, and make sure the governance framework, operational process and vendor contract each provide a defensible answer before a regulator, customer or claimant asks the question first.

Join ACC for more AI insights!

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

Matías Cañibe

Associate General Counsel and Data Protection Officer based

Luxembourg

Matias Cañibe is Associate General Counsel, Assistant Corporate Secretary and Data Protection Officer at a NASDAQ-listed multinational company. He advises Boards, executive management and

business teams on corporate governance, privacy, cybersecurity, artificial intelligence governance and cross-border legal matters. He holds law degrees from two civil-law jurisdictions and an MBA. He writes on the legal and governance implications of emerging technologies.