

ARTICLE

The Future-Proof Fantasy of AI Regulation

Thibault Schrepel

The Network Law Review is pleased to present a special issue entitled “The Law & Technology & Economics of AI.” This issue brings together multiple disciplines around a central question: What kind of governance does AI demand? A workshop with all the contributors took place on May 22–23, 2025, in Hong Kong, hosted by Adrian Kuenzler (HKU Law School), Thibault Schrepel (Vrije Universiteit Amsterdam), and Volker Stocker (Weizenbaum Institute). They also serve as the editors.

Abstract: The EU’s quest for “future-proof” AI regulation is a fantasy. AI evolves through emergent properties that defy prediction, yet Brussels continues to draft rules with an industrial, linear mindset. The result is a regulatory immune system that can detect but not respond. The path forward is adaptive regulation: modular rules, real-time sensing, plural triggers, and institutional memory.

1. Introduction

The regulatory challenge posed by artificial intelligence exposes a fundamental mismatch between our legal frameworks and the systems they seek to govern. While policymakers globally rush to craft “future-proof” AI regulations (i.e., rules designed to endure technological change unchanged), this approach is structurally inadequate for governing complex adaptive systems like AI.

Consider the physics: water obeys thermodynamic laws that remain constant, making waterproof materials possible. Fire obeys the kinetics of combustion, making fireproof materials possible. But technology evolves in ways that violate our regulatory assumptions at their core. Technological evolution does not simply

accelerate existing patterns but fundamentally rewrites them. Each generation of AI does not just improve, it recombines and spawns unexpected capabilities. GPT-1 to GPT-4 was not iteration; it was metamorphosis.

EU digital regulation assume they can capture these phase transitions in advance. But complexity science teaches us otherwise. Emergent properties cannot be anticipated from initial conditions. The Commission's scramble to insert a general-purpose AI model chapter into the AI Act mid-drafting proves the point: emergence defeats prediction.

Against this background, I analyzed whether EU digital regulation compensates for prediction failure through adaptation mechanisms. The dataset: eight Digital Acts (DGA, DMA, DSA, DORA, Chips Act, Data Act, AI Act, CRA) adopted between 2022-2024 and evaluated against fourteen adaptivity criteria. The finding: these Acts contain adaptive machinery such as review clauses, delegated acts, and monitoring obligations but lack adaptive capacity. They built sensors without reflexes, feedback loops without response functions. This short piece distills the key insights for AI governance. The longer piece provides more details, data, and extend the scope of analyze.¹

2. The current state of EU regulation

AI systems exhibit the hallmarks of complexity. They show non-linear dynamics where small changes produce large effects, emergent properties not deducible from components, and continuous adaptation rather than equilibrium states. When a large language model develops unexpected capabilities through scaling, or when reinforcement learning systems discover novel strategies their creators never anticipated, we witness complexity in action. Regulating such systems with tools designed for predictable, linear phenomena is like using Newtonian mechanics to describe quantum behavior. The regulatory toolkit assumes Gaussian distributions and mean reversion. AI exhibits power laws and runaway effects.

Meanwhile, AI regulation approach AI systems as a static technology. This is a trend across EU Digital Acts; they remain fundamentally frozen at their moment of enactment. Yes, they incorporate some adaptation mechanisms, including review clauses, delegated acts, and monitoring obligations, but they remain anchored in neoclassical assumptions about predictable technological trajectories.

The numbers expose the fiction. Start with the sensing deficit. About 75% of the EU Digital Acts mandate some form of data collection as to the effects they produce, but only the DSA requires real-time monitoring—and even then, only for very large

¹ Thibault Schrepel, Adaptive Regulation (Amsterdam Law & Technology Institute, Aug. 15, 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5416454

platforms. Machine-readable reporting, essential for computational oversight, appears in just two of eight Acts. Most tellingly, adaptation triggers remain almost entirely discretionary. Six Acts deploy escape clauses (“where appropriate,” “where necessary”) that convert mandatory reviews into optional exercises rather than defining concrete thresholds that would mandate review. The Commission holds monopolistic control over initiating adaptations across all eight Acts, which creates a single point of failure where distributed intelligence is most needed. Review cycles, when they exist, typically stretch to three or five years; geological time in digital markets where capabilities can transform in weeks. In complexity science terms, EU Digital Acts lack real adaptive capacity, the defining characteristic that separates living systems from dead matter. Put differently, the EU built an immune system that detects pathogens but cannot produce antibodies.

The AI Act illustrates the limits of adaptive capacity no less than others. It multiplies review clauses (Article 112) and invents institutions, the AI Office, the AI Board, the Advisory Forum. They are bound to a single center. The Commission alone decides. No national authority, no independent agency can pull the trigger. Review cycles are annual for prohibited and high-risk AI, quadrennial for governance. In digital time, both are glacial. Machine-readable reporting is absent. Oversight clings to prose and PDFs, not computable data. Exogenous triggers are listed (including technological breakthroughs, evolving risks to health, safety or fundamental rights, and shifts in the information society) but the Act frames them in discretionary terms. Article 112(10) provides that the Commission may propose revisions on these grounds. It does not say it shall. The difference is decisive. What looks like an automatic reflex is in fact an option left to political convenience. And when adaptation comes, it often moves in only one direction: Article 7 empowers the Commission to add new high-risk uses to Annex III, but not to remove them.

3. How to make EU regulation adaptive

The relevance of EU regulation depends on its ability to function as a complex adaptive system, structured around four principles. These apply beyond the AI focus of this short piece.

First, modular regulatory architecture that separates fundamental principles (which remain stable) from operational rules (which adapt rapidly). For AI, modular regulatory architecture means keeping liability frameworks and fundamental rights protections in primary law while allowing technical thresholds (model size triggers, compute limits, risk assessment methodologies) to evolve through secondary instruments. This also means a list of what are considered essential and non-essential elements that can be tackled by delegated and implementing acts.

Today's EU Digital Acts show limited modularity. The AI Act freezes its "AI system" definition in Article 3 while allowing only annexes to be modified through delegated acts. The CRA can only modify product categories in Annexes III and IV, not core definitions. Six Acts technically allow any modification, but only through full legislative procedures taking years. The result: technical standards and risk categories remain as frozen as constitutional principles.

Second, distributed sensing that monitors AI systems continuously rather than episodically. The DSA's API-based oversight of platforms offers a template. Imagine similar real-time monitoring for foundation models, with machine-readable incident reports and performance metrics flowing continuously to regulators.

For now, the empirical evidence reveals a troubling absence of machine-readable requirements outside the DSA. The Data Governance Act (ostensibly about data governance) contains no machine-readable reporting obligations. The Cyber Resilience Act collects compliance data in prose, like asking for source code via fax. This analog approach to digital oversight guarantees that regulatory learning will lag technological development by design, not accident. When regulations cannot computationally process their own performance data, adaptation becomes aspiration rather than operation.

These Digital Acts also collect data in annual batches (at best) when AI capabilities can transform in weeks. They are taking yearly photographs of a Formula One race. DORA requires incident reports within 72 hours (Article 19.4), which sounds responsive until you realize that is enough time for an AI model to be deployed globally, cause harm, and spawn three competitor versions. The CRA's first review is not until 2030 (Article 70.1). By then, today's AI will look like cave paintings.

Third, pluralistic triggering mechanisms that allow multiple actors to initiate regulatory adaptation. National authorities observing local AI harms, technical bodies detecting capability jumps, or civil society documenting systematic biases should all possess formal powers to trigger review processes when predetermined thresholds are met.

As of today, the Chips Act and AI Act stand alone in defining concrete performance indicators (Annex II of the Chips Act specifies metrics like SME participation rates, infrastructure access rates, and venture capital flows; the AI Act's Article 112.11 mandates the AI Office to develop a risk-based evaluation methodology). The other six Acts rely on vague "impact assessment" language. The DMA's Article 53 simply calls for evaluation "where appropriate," regulatory speak for "maybe never." The DGA's Article 35 mentions assessing "impact" without metrics, and the Data Act's Article 49 lists review topics but no measurable triggers.

Regarding who can initiate revisions when thresholds are met, all eight Digital Acts grant this power exclusively to the European Commission. Member States, national agencies, and sectoral bodies remain advisory at best. This monopoly on adaptation initiative creates a single point of failure precisely where distributed intelligence is most needed. Even the relatively advanced AI Act, with its multi-tiered governance structure, cannot escape this centralized bottleneck.

Fourth, networked institutional memory that connects learning across regulatory domains. AI governance cannot occur in isolation. Changes in data protection, competition, or product safety rules cascade through AI systems. Cross-regulatory coordination bodies must ensure adaptations remain coherent across the regulatory ecosystem.

The AI Act’s creation of the AI Office, Board, and Advisory Forum represents the most sophisticated institutional learning architecture among the Acts studied. The Board can propose amendments (Article 66.e.vii) and the Advisory Forum produces written contributions (Article 67.8). Yet even this falls short of true networked memory. None of these Acts creates cross-regulatory coordination mechanisms. They leave each regulation to evolve in isolation despite governing interconnected digital markets.

4. The road ahead, and the what if

Applying these principles to the EU regulatory framework – starting with AI regulation – would constitute nothing short of a (much needed) revolution in digital governance.² It will not be easily done. Perhaps the most radical implication of adaptive regulation is accepting that specific rules must die for regulatory systems to live. Just as biological systems achieve resilience through cellular turnover, regulatory frameworks need mechanisms for controlled obsolescence and renewal.

The alternative to adaptive regulation is not stable regulation, it is ossified regulation. Static rules for dynamic systems create three pathologies. They become either irrelevant (as technology routes around them), harmful (as they lock in outdated assumptions), or both (as they provide false certainty while failing to address real risks). The recent emergence of agentic AI systems illustrates this perfectly. Regulations written for supervised learning models now confront systems that act ‘autonomously’ and interact with other AI agents in ways that produce emergent collective behaviors. No amount of definitional elasticity in original legislative text could have anticipated these developments.

The choice is not between regulatory permanence and regulatory volatility. It is between regulations that assume technological stasis and those that factor in

² Do not worry, I realize how cliché it is for a Frenchman to call for a revolution.

continuous change. It is between future proof and future responsive regulation. As AI systems grow more powerful and more pervasive, our regulatory frameworks must evolve from monuments to living systems. Continuing to pretend our regulation can be “future proof” or adapt already is not good policymaking, it is fantasy.

Thibault Schrepel