Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

European Journal of Risk Regulation (2026), 1–14
doi:10.1017/err.2026.10165

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

From the Law of the Horse to the Workflow: Why Institutions, Not the Model, Decide the EU AI Act's Effectiveness: Lessons from Hungary's First-Wave Implementation

Nimród Mike

Institute of Data Analytics and Information Systems, Corvinus University of Budapest, Hungary
Email: nimrod.mike@uni-corvinus.hu

Abstract

Whether a new technology needs its own law is an old question, and for artificial intelligence (AI), it has been posed at the wrong level of abstraction. Both the case against technology-specific law and the case for regulation by design treat the AI model as the regulated object, although a general-purpose model acquires its risk only when deployed. This article relocates the regulatory object to the deployment workflow, defined as the envelope of thresholds, escalation pathways and human checkpoints that surrounds a model in use, and reads the EU Artificial Intelligence Act as management-based regulation that works only where the workflow is specifiable and a standing supervisor examines it. Hungary's first-wave implementation shows the second condition deciding the outcome: the same obligations are supervised by the central bank in finance and by a ministry with no examination relationship in employment, so the Act's uniform text conceals uneven supervision within one Member State. The Digital Omnibus on AI, which postpones the high-risk obligations because standards and authorities are late, opens a window to build supervision deliberately: the priority is not better definitions but better supervision.

Keywords: EU AI Act; management-based regulation; national competent authorities

I. Introduction

Whether artificial intelligence (AI) needs a law of its own is an old question in new dress, and it descends from a quarrel about cyberspace. In 1996 Judge Frank Easterbrook warned against a "law of the horse," a body of doctrine assembled around a single technology. Such a field, he argued, is "doomed to be shallow and to miss unifying principles," and the general fields of property, tort and contract offer the surer route to governing novelty.[1] Three years later Lawrence Lessig replied that studying a technology can illuminate the general law, because cyberspace reveals a regulatory force that orthodox theory neglects: architecture, in which digital environments takes the form of code and constrains conduct as effectively as any rule.[2] The two positions still frame the debate. Easterbrook counsels against a dedicated statute and in favour of the general law that already exists, while Lessig supplies the intellectual licence for the European Union's Artificial Intelligence Act

[1] F. H. Easterbrook, "Cyberspace and the Law of the Horse" (1996) University of Chicago Legal Forum 207, 207–8.
[2] L. Lessig, "The Law of the Horse: What Cyberlaw Might Teach" (1999) 113 Harvard Law Review 501; see also L. Lessig, *Code: Version 2.0* (Basic Books 2006) pp 121–5.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

(AI Act),[3] whose duties on data governance, documentation, risk management and human oversight read as design mandates rather than as prohibitions.

This article argues that the debate has been conducted at the wrong level of abstraction, and that once the level of analysis is corrected, the disagreement dissolves rather than being settled in either side's favour. When the two positions are transposed to AI, both come to rest on a premise that neither Easterbrook nor Lessig ever advanced: that the object of regulation is the AI model itself. That premise fails most clearly for the general-purpose, generative models now at the centre of the policy debate. Such a model is built once and reused across countless applications, and it acquires a determinate purpose, and therefore a determinate risk, only when it is put to work in a particular deployment. This does not mean that the model cannot be regulated. It means that the model cannot be the object of the Act's risk-based regime, because at the level of the model there is no determinate risk to classify and no legible architecture to audit. The productive unit of analysis is instead the deployment workflow: the system that wraps a probabilistic model in a deterministic envelope of decision rules, confidence thresholds, escalation pathways, override logic and human checkpoints. At that level Lessig's insight returns, because the envelope constrains what the system may do and because it can be specified, documented and audited; and Easterbrook's insight returns too, because whether the envelope binds in practice depends on the ordinary supervisory institutions behind it.

Read in this way, the Act is an instance of what regulatory theory calls management-based regulation, a technique that governs the management process around a product rather than the product or the technology itself.[4] The technique's record is conditional in a specific way: because the firm designs its own controls, the technique works only where a regulator has the capacity and, above all, the continuing relationship with the regulated firm that evaluating those controls requires.[5] We distil this into two conditions for the Act's high-risk regime to achieve regulatory sufficiency, by which we mean that the regime actually contains the risk it addresses. First, the workflow envelope must be specifiable enough to absorb the model's residual risk. Second, a supervisory authority with a continuing examination relationship must exist to convert paper obligations into accountability, that is, into obligations that change conduct rather than only documentation.

We test the argument against the first-wave implementation of the Act in Hungary. The Act routes the supervision of AI used by financial institutions to the financial supervisor and of AI in medical devices to the medical-device authority, but for standalone high-risk domains such as employment it leaves the choice of authority to each Member State,[6] and Hungary chose a government ministry. The result is close to a controlled comparison: supervision of identical obligations falls to the central bank in finance, to the medical-device authority in healthcare, and to a ministry with no examination relationship in employment. A change of government in May 2026 and the Digital Omnibus on AI of July 2026, which postpones the high-risk obligations to December 2027 and August 2028, have reopened the design of that supervision.[7] Section II relocates the regulatory object. Section III reads the Act as management-based regulation and shows the Omnibus conceding the

[3] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L, 2024/1689 ("AI Act").

[4] C. Coglianese and D. Lazer, "Management-Based Regulation: Prescribing Private Management to Achieve Public Goals" (2003) 37 Law & Society Review 691.

[5] J. Black and R. Baldwin, "Really Responsive Risk-Based Regulation" (2010) 32 Law & Policy 181; S. Gilad, "It Runs in the Family: Meta-Regulation and Its Siblings" (2010) 4 Regulation & Governance 485.

[6] AI Act, *supra*, note 3, Arts. 70 and 74.

[7] Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI) [2026] OJ L, 2026/1744 ("Digital Omnibus").

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

Act's two gaps. Section IV develops the Hungarian comparison. Section V argues that the postponement is a window in which to replace a placeholder with deliberate architecture.

II. Relocating the regulatory object

The two positions are usually read as opposites, but they have more in common than that reading allows. Easterbrook's scepticism rests on a claim about the object: general doctrine can absorb a new technology because the technology is tractable, a thing to which settled principles attach. Lessig accepts the goal of illuminating the general law and adds a fourth modality of regulation, architecture, alongside law, social norms and markets. In digital environments, he argues, architecture is code, and because code constrains behaviour directly, a regulator can govern by prescribing the architecture itself, that is, by mandating how the code is designed. The European legislator took the Lessigian path, though not exclusively. Article 5 of the Act forbids certain practices outright, in the classical form of a conduct rule; but the bulk of the text, the high-risk regime, prescribes how systems must be built, trained, documented and overseen, and it is this larger part that follows Lessig.

The architectural reading faces a difficulty that its proponents rarely confront, and the difficulty is most acute for the general-purpose models that now dominate the debate. Lessig's "code is law" assumes that code is deliberately designed and that its rules can be read back from the artefact. A general-purpose, generative model does not fit that picture: its behaviour is not written by its designers but induced from data through a training process that tunes billions of parameters. The result, in the words of Judge, Nitzberg and Russell, is a system for which "code is no longer law," since the program text neither determines how the model behaves nor makes that behaviour legible,[8] so that the decision function cannot be audited line by line against a legal standard as a conventional program can.[9] If the model is the regulated object, then Lessig's insight fails, because the architecture is exactly what cannot be read. Easterbrook fares no better, because general principles require an object stable enough to be subsumed under them, and an emergent function that can shift after deployment offers no such stability. At the level of the model the debate reaches an impasse.

The impasse is an artefact of the chosen unit of analysis. A general-purpose model is not deployed as such. It is built once and then implemented in many applications, so that the same model may sit behind a marketing assistant, a loan adjudicator and a clinical triage tool, carrying a different risk in each because it has no fixed purpose or risk of its own. This is not to say that the model escapes regulation. A model is software and, as such, a product within the reach of general product-liability law, and the Act adds obligations addressed to the providers of general-purpose models, concerning documentation, training-data summaries, copyright policy and, for models with systemic risk, evaluation and incident reporting.[10] Those duties are notable for their form: they govern the provider's development process rather than the model's behaviour, as one would expect where behaviour cannot be read. What can be regulated at the level of the model is the conduct of its maker; what cannot be regulated there is the risk, because the risk does not yet exist. The question Judge, Nitzberg and Russell leave open, how to govern a system whose behaviour cannot be read off its code, is answered at the level of implementation.

[8] B. Judge, M. Nitzberg and S. Russell, "When Code Isn't Law: Rethinking Regulation for Artificial Intelligence" (2025) 44 Policy and Society 85.

[9] J Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms" (2016) 3 Big Data & Society 1.

[10] AI Act, *supra*, note 3, Arts. 53 and 55. On software as a product under the revised product-liability regime, see infra, note 17.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

That implementation is the deployment workflow. A deployed system is never only a model; it is a model inside a workflow that decides what the model may determine, when its output is acted on automatically, when a case is escalated, and where a human must step in. This envelope is deterministic: it is designed, legible and stable. At this level Lessig's thesis returns, because code is law again, not in the model, whose parameters no one can read, but in the surrounding architecture that bounds what the model is allowed to do. The envelope can be specified in advance, documented and audited against a legal standard even though the model inside it cannot; at the workflow level, the model's opacity becomes a property of one component whose outputs the surrounding controls are built to contain.

The concept needs precise edges, because the Act comes close to it without naming it. The Act's "intended purpose" fixes the use for which the provider markets a system, and so names the workflow's goal rather than its controls.[11] The deployer duties of Article 26 assign obligations to a person and presuppose an envelope that someone has already designed. The sociotechnical turn in the algorithmic–accountability literature, moreover, insists that a model be judged within its social context.[12] Our claim is narrower and specifically legal. The workflow is an object, not an actor: it is the engineered part of that context, built by the provider, configured by the deployer, and legible to an auditor in a way that neither the model's parameters nor the social setting can be. The Act allocates duties to actors and gestures at the object between them, but never makes it the unit of regulation. This article does, because the workflow is the one artefact in a deployment that can be specified in advance, inspected in place, and held constant across the actors responsible for it.

This move also names the technique the Act uses. Regulatory theory distinguishes three families: technology-based regulation prescribes a means, performance-based regulation prescribes an outcome, and management-based regulation requires the firm to design and run internal processes aimed at a public goal.[13] The Act's high-risk regime belongs to the third family, since its duties are process duties directed at the workflow rather than specifications of the model or guarantees of its output. The naming matters because the obvious objection, that all regulation depends on enforcement, misses what is distinctive about this technique. A technology-based or performance-based rule states its own standard and can be verified in a single inspection. A management-based rule devolves the design of the controls to the firm, so that the regulator must judge the adequacy of a self-designed process which the firm alone sees in full. That judgement cannot be made episodically; it requires access, accumulated familiarity and the possibility of return. The evidence is consistent: the technique delivers where the regulator has capacity, a stable mandate and a continuing relationship with the regulated, and underperforms elsewhere.[14] The same failure mode appears inside the firm: where no external authority is placed to test them, internal compliance processes drift, and Selbst describes the drift as a managerial move in which substantive questions about social impact become questions about reputational risk and documentation grows symbolically adequate but substantively thin.[15] This is not an argument against process obligations but an argument that they are inert unless a supervisor stands ready to interrogate them: the same statutory duties produce accountability where a capable supervisor examines them and paperwork where none does.

[11] AI Act, *supra*, note 3, Arts. 3(12), 3(23), 25 and 26.

[12] A. D. Selbst and others, "Fairness and Abstraction in Sociotechnical Systems" in *Proceedings of the Conference on Fairness, Accountability, and Transparency* (ACM 2019) 59.

[13] Coglianese and Lazer, *supra*, note 4, pp 691–3.

[14] Gilad, *supra*, note 5; N. Gunningham and D. Sinclair, "Organizational Trust and the Limits of Management-Based Regulation" (2009) 43 Law & Society Review 865.

[15] A. D. Selbst, "An Institutional View of Algorithmic Impact Assessments" (2021) 35 Harvard Journal of Law & Technology 117.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

Two conditions therefore decide whether the Act's high-risk regime achieves regulatory sufficiency. The first is specifiability: the workflow envelope must be pinned down tightly enough to absorb the model's residual error within tolerance. Where the output is bounded, as when a transaction is fraudulent or not, or an application is approved or declined, the envelope can be specified; where the output is open-ended, the envelope cannot hold and the residual risk escapes. The second is supervisory backing: a standing authority with an examination relationship must exist to test whether the mandated processes are real, since without one, management-based obligations are discharged on paper. The Act's own catalogue of process duties in Articles 9 to 17 satisfies the first condition only partly, and the Act satisfies the second not at all, because it leaves supervision to the Member States.[16]

We present the two conditions as necessary rather than sufficient, and other mechanisms of accountability are conceivable, so the alternatives deserve an answer. Deterrent fines, private claims and *ex post* liability might each, in principle, hold a deployer to account without a standing supervisor, but none of them does. Fines presuppose an authority with the standing to find the breach, and so fold into the second condition. Private enforcement under the Act is thin, and the information asymmetry that defeats an external auditor weighs even more heavily on an individual claimant. Liability has been left to general law, since in 2025 the Commission withdrew its proposal for an AI Liability Directive, an Easterbrookian choice in itself, in favour of the revised Product Liability Directive, which now reaches software and AI systems;[17] but product liability compensates damage caused by defective products, not decisional harms. What remains is the front-end relationship that the second condition describes.

III. The AI Act as management-based regulation

Read through this lens, the high-risk regime is plainly management-based. Article 9 requires a continuous risk-management system across the lifecycle; Article 10 governs the data pipelines that feed the model; Articles 11 and 12 require technical documentation and record-keeping; Article 14 requires that the system be designed for effective human oversight; and Article 17 requires a quality-management system over the whole. For most standalone systems the provider certifies conformity by its own assessment against these process requirements and signals it with a CE marking borrowed from the Union's product-safety machinery.[18] None of these provisions regulates the model as such; each regulates the workflow around it. At its core, the Act is a management-based instrument in the dress of product safety.

That borrowing has consequences, as Edwards observes: the Act is held in place by the New Legislative Framework (NLF) and its internal-market legal basis, which cast it as a product-safety instrument and entrust its enforcement to market surveillance authorities, a class of regulator historically less interventionist than the data-protection authorities that drove enforcement under the General Data Protection Regulation (GDPR).[19] Almada and Petit put the structural point sharply: the Act joins two traditions of EU law, product safety and fundamental-rights protection, that do not naturally cohere, and the seam

[16] AI Act, *supra*, note 3, Arts. 9–17.

[17] C. Novelli, P. Hacker, J. Morley, J. Trondal and L. Floridi, "A Robust Governance for the AI Act: AI Office, AI Board, Scientific Panel, and National Authorities" (2025) 16 European Journal of Risk Regulation 566. On the withdrawal of the AI Liability Directive proposal, see European Commission, Commission Work Programme 2025, COM(2025) 45 final, Annex IV. Directive (EU) 2024/2853 on liability for defective products [2024] OJ L, 2024/2853, extends the product-liability regime to software, including AI systems.

[18] AI Act, *supra*, note 3, Arts. 43 and 48; Annexes VI and VII.

[19] L. Edwards, "Regulating AI in Europe: Four Problems and Four Solutions" (Ada Lovelace Institute 2022).

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

shows in three ways, a pacing problem, a mismatch between the regulatory lens and its object, and a path dependence on standardisation bodies ill-suited to questions of rights.[20] On our account each of these is a symptom of a single choice, that of regulating a workflow through machinery built for a product. Hungarian scholarship reads the Act the same way, as the newest entry in the NLF,[21] and since the NLF's tools are conformity assessment, documentation and presumptions of conformity, the recurrence of that reading across jurisdictions is itself a sign that the Act's centre of gravity lies in process and documentation, that is, in the workflow.

Reading the Act as management-based regulation also dissolves a long controversy over the Article 3(1) definition of an "AI system" and the weight placed on the verb "infers."[22] If the regulated object were the model, the definition would have to separate the systems that warrant the full high-risk regime from ordinary software, and critics are right that it cannot: it is at once over- and under-inclusive.[23] Ebers adds the mirror-image point that the Act, applying its heaviest duties by category rather than by assessed risk, can capture deterministic, well-understood software.[24] Both objections have force at the level of the model and fade at the level of the workflow, because what must be specified and supervised is the envelope of decision governance, and that envelope is specifiable whether or not the component at its centre is best called "artificial intelligence." The definition fixes the scope of the obligations; it does not decide whether they can be made to work.

The same reframing explains the shape of the general-purpose AI provisions, the clearest doctrinal symptom of the problem we identify. Articles 51 to 55 regulate general-purpose models at the level of the model, by compute thresholds and model-level documentation, and the legislator had little choice, since such models have no determinate risk until embedded in a downstream application.[25] That a separate, model-level regime was needed is itself evidence for the point. Model-level classification cannot carry the Act's risk logic, since the same model may power a trivial productivity tool or a clinical decision aid, and only the workflow determines which of these it is. Tellingly, the regime built for the model consists of process duties on the provider rather than behavioural requirements on the model. The general-purpose chapter is the exception that proves the rule.

If the Act is management-based regulation, its two weaknesses follow from the two conditions. The first is that it under-specifies the envelope. Articles 11 and 17 require "technical documentation" and a "quality management system," but they do not say what an auditable workflow looks like: the decision thresholds, escalation criteria, override logic and measures of oversight effectiveness that would make the envelope independently verifiable.[26] These obligations sit at too high a level of abstraction to be tested, and because most standalone systems are self-assessed, an under-specified duty leaves providers wide latitude to declare their own systems compliant.[27]

[20] M. Almada and N. Petit, "The EU AI Act: Between the Rock of Product Safety and the Hard Place of Fundamental Rights" (2025) 62 Common Market Law Review 85.

[21] A. Tóth, "Az Európai Unió Mesterséges Intelligencia Törvényéről" [On the European Union's Artificial Intelligence Act] (2024) *Gazdaság és Jog* 2024/5–6, 3; K. Mezei, "A mesterséges intelligencia jogi szabályozásának aktuális kérdései az Európai Unióban" [Current Issues in the Regulation of Artificial Intelligence in the European Union] (2023) *In Medias Res* 2023/1, 53.

[22] AI Act, *supra*, note 3, Art. 3(1).

[23] M. Ebers, "Truly Risk-Based Regulation of Artificial Intelligence: How to Implement the EU's AI Act" (2025) 16 European Journal of Risk Regulation 684; M. Veale and F. Zuiderveen Borgesius, "Demystifying the Draft EU Artificial Intelligence Act" (2021) 22 Computer Law Review International 97.

[24] Ebers, *supra*, note 23, pp. 690–5.

[25] AI Act, *supra*, note 3, Arts. 51–5.

[26] AI Act, *supra*, note 3, Arts. 11 and 17, and Annex IV.

[27] Veale and Zuiderveen Borgesius, *supra*, note 23, on the latitude that internal self-assessment leaves to providers.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

The second weakness is graver. The Act does not supply the supervisory relationship that management-based regulation needs; it delegates that relationship to the Member States. Article 74 routes market surveillance of AI in regulated products, including medical devices, to the authority named under the relevant sectoral law, and market surveillance of AI used by financial institutions to the financial supervisor. For standalone high-risk systems in Annex III outside the law-enforcement, migration, justice and democratic-process categories, however, it leaves the identity, and with it the capacity, of the market surveillance authority to each Member State's designation under Article 70.[28] The Act therefore presupposes a supervisor that it does not itself create, and whether its obligations are enforced in practice depends on an institutional choice, the designation of the national authority, that the Act leaves to the Member States.

Both weaknesses have now surfaced in the Act's first major amendment. The Digital Omnibus on AI, published in the Official Journal in July 2026, resets the application of the high-risk regime to 2 December 2027 for standalone Annex III systems and to 2 August 2028 for product-embedded ones.[29] The recitals justify the delay by "the delayed availability of standards, common specifications, and alternative guidance and the delayed establishment of national competent authorities."[30] These are our two conditions, described as unmet in the legislator's own words: the envelope lacks operational content because the standards are late, and the supervisory relationship is missing because most Member States have not built it. The postponement amounts to the Union conceding the diagnosis.

The Omnibus does more than delay; it also redraws part of the supervisory map. The amended Article 75 makes the Commission, acting through the AI Office, exclusively competent to supervise two classes of system: those built on a general-purpose model and supplied by the model's own provider or a provider within the same undertaking, and those that constitute, or are embedded in, a very large online platform or very large online search engine designated under the Digital Services Act.[31] Exclusive competence follows the provider and reaches deployers only where they belong to the same undertaking; every other deployer remains under national supervision.[32] The Union has built a standing supervisor exactly where the workflow sits in the provider's own hands, because that is the only workflow a central authority can reach; where model and workflow part company, as they ordinarily do in employment, healthcare and finance, supervision still turns on the national designation. The centralisation cures the supervisory gap for a small class of vertically integrated providers and confirms it for everyone else.

IV. The Hungarian case

1. The designation

Hungary implemented the Act through Act LXXV of 2025 and an implementing decree, Government Decree 344/2025 (X.31.).[33] The National Accreditation Authority became the notifying authority, while the minister responsible for enterprise development is both the

[28] AI Act, *supra*, note 3, Arts. 74(3), 74(6) and 74(8); standalone high-risk systems otherwise fall to the authority designated under Art. 70.

[29] Digital Omnibus, *supra*, note 7, amending Art. 113, third paragraph, of the AI Act.

[30] Digital Omnibus, *supra*, note 7, recital 40.

[31] AI Act, *supra*, note 3, Art. 75(1), as replaced by the Digital Omnibus, *supra*, note 7, and, on the AI Office's powers, Arts. 75a–75d, as inserted by it; Regulation (EU) 2022/2065 (Digital Services Act) [2022] OJ L277/1, Art. 33.

[32] Digital Omnibus, *supra*, note 7, recital 31: "Other deployers should remain subject to national supervision and enforcement."

[33] Act LXXV of 2025 on the implementation in Hungary of Regulation (EU) 2024/1689; Government Decree 344/2025 (X.31.) on its implementation. See L. Pók, "A mesterséges intelligencia rendszerek hatósági felügyelete Magyarországon: máris változhat?" [Regulatory Supervision of AI Systems in Hungary: About to Change Already?] (GDPR blog, 12 May 2026) <https://gdpr.blog.hu> (last accessed 14 June 2026).

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

market surveillance authority under Article 70 and the single point of contact, with the tasks discharged by a dedicated market-surveillance department inside the ministry. Financial-sector AI is the exception: following the Act's own allocation, Act LXXV assigns market surveillance to the Magyar Nemzeti Bank (MNB) where a system is directly connected to the provision of financial services.[34] An advisory Hungarian Artificial Intelligence Council completes the framework, drawing together the main sectoral regulators and academic and industry bodies to coordinate the national authorities' AI activities, but holding no decision-making power.[35]

The horizontal choice is notable. Across most of the economy, the default supervisor of high-risk AI is a government ministry, although the Act requires national competent authorities to exercise their powers "independently, impartially and without bias," in line with Regulation (EU) 2019/1020 on market surveillance.[36] A department inside a ministry that also promotes the digital economy is not the natural vehicle for that independence, and among Member States a ministerial designation is a rarity.[37] The designation also departs from Hungary's own plan, since the national AI Strategy 2025–2030 had envisaged a dedicated, independent Hungarian Artificial Intelligence Office rather than a ministerial department.[38]

The independence requirement has more content than its brevity suggests, because of what the systems in question do. High-risk systems in employment, credit scoring and healthcare almost by definition process personal data, and their operation therefore constitutes processing within the meaning of the GDPR. Union law has long required that the supervision of such processing be entrusted to an authority acting "with complete independence." That requirement has constitutional rank under the Charter and the Treaties, and the Court of Justice has enforced it strictly: it excludes any direct or indirect external influence, and in proceedings against Hungary the Court held that the premature termination of the data-protection supervisor's term of office infringed it.[39] The Act extends that standard to the categories of high-risk system it regards as most sensitive: for law-enforcement, migration, justice and democratic-process systems, Article 74(8) requires Member States to designate as market surveillance authority either the data-protection authority or a body established under the same conditions of independence.[40] The European Data Protection Board has recommended that Member States consider the data-protection authority for the remaining Annex III systems as well, because of that independence.[41] Its premise is sound: the Union already knows what independent supervision of automated decisions about persons looks like, and it does not look like a ministerial department.

One caveat is in order before the comparison. Because the high-risk obligations do not yet apply, no enforcement record exists in any sector, and the comparison therefore concerns structural preconditions rather than outcomes: the three sectors face the same

[34] AI Act, *supra*, note 3, Art. 74(6); Act LXXV of 2025, *supra*, note 33, s 5(1).

[35] Act LXXV of 2025, *supra*, note 33; the advisory council's coordinating function draws on AI Act, *supra*, note 3, Art. 74(10).

[36] AI Act, *supra*, note 3, Art. 70(1); Regulation (EU) 2019/1020 on market surveillance and compliance of products [2019] OJ L169/1, Art. 11.

[37] Pók, *supra*, note 33; European Commission, "Market Surveillance Authorities under the AI Act" <https://digital-strategy.ec.europa.eu/en/policies/market-surveillance-authorities-under-ai-act> (last accessed 29 August 2026).

[38] Hungary's Artificial Intelligence Strategy 2025–2030, s 5.1.

[39] Regulation (EU) 2016/679 (General Data Protection Regulation) [2016] OJ L119/1 ("GDPR"), Art. 52(1); Charter of Fundamental Rights of the European Union, Art. 8(3); Art. 16(2) TFEU; Case C-518/07 *Commission v Germany* ECLI:EU:C:2010:125; Case C-614/10 *Commission v Austria* ECLI:EU:C:2012:631; Case C-288/12 *Commission v Hungary* ECLI:EU:C:2014:237.

[40] AI Act, *supra*, note 3, Art. 74(8), referring to the conditions laid down in Arts. 41–4 of Directive (EU) 2016/680.

[41] European Data Protection Board, "Statement 3/2024 on data protection authorities' role in the Artificial Intelligence Act framework" (16 July 2024) paras 8–9.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

statutory duties but differ sharply in the institutional capacity available to give those duties practical effect. The evidence lies in the supervisors' existing practice, not in results produced under the Act itself.

2. Financial services: both conditions present

Financial services is the sector where both conditions hold. The high-risk use cases are creditworthiness assessment and credit scoring;[42] fraud detection and anti-money-laundering triage, though excluded from the high-risk list, share the same bounded-output structure. Because the outputs are bounded, the workflow envelope can be specified with precision, as scorecards, decision thresholds, override and exception handling, and human review of adverse decisions.

The second condition is met because the supervisor already exists and has been designated. The MNB supervises credit institutions as a matter of routine, examining their internal-control arrangements through the periodic supervisory review that prudential regulation already requires, and in 2025 it began to extend that supervision to AI, recommending that institutions govern their use of AI within existing risk-management structures.[43] The contrast with a one-off conformity assessment is decisive. A supervisor in a standing relationship can return to a firm, re-examine its controls and test whether the documented workflow is the operating one. What the supervisor audits is not the internals of any model but the governance architecture around it, which is precisely the object that management-based regulation renders auditable. The Act's classification of credit-scoring AI as high-risk does not so much create a new burden as formalise an existing one; here the technique works, because the institution it assumes is in place.

3. Healthcare: object-specific supervision, partial specifiability

Healthcare tests the first condition while the second holds in a qualified form. AI that is, or is a safety component of, a medical device requiring third-party conformity assessment is high-risk through the Act's interface with the Medical Devices Regulation,[44] and Article 74(3) routes its market surveillance to the authority named under that Regulation, which in Hungary is the medical-device competent authority.[45] A sectoral supervisor with an object-specific relationship therefore exists. Its reach, however, ends at the artefact. Notified bodies certify the device and post-market surveillance follows the product, but neither examines how a deploying hospital sets its thresholds, staffs its review or acts on the outputs. Supervision covers the workflow's components, not the workflow as operated, even though clinical practice already distributes post-market duties across the chain of manufacturer, deploying institution and clinician.[46]

This arrangement works for bounded clinical AI: image triage, diagnostic prioritisation, sepsis scoring and similar classification tasks with reasonably well-defined ground truth. For these the workflow envelope is specifiable as the set of conditions under which an output is acted on, flagged or referred to a clinician, and the conformity-assessment and

[42] AI Act, *supra*, note 3, Annex III, point 5(b).

[43] Magyar Nemzeti Bank, Recommendation 13/2025 (XII.3.) on the digital transformation of credit institutions; AI Act, *supra*, note 3, Art. 74(6).

[44] AI Act, *supra*, note 3, Art. 6(1) and Annex I; Regulation (EU) 2017/745 on medical devices [2017] OJ L117/1 ("MDR").

[45] AI Act, *supra*, note 3, Art. 74(3). In Hungary the National Centre for Public Health and Pharmacy (NNGYK), successor since 2023 to the National Institute of Pharmacy and Nutrition (OGYÉI), is both the competent authority for medical devices and the designating authority for notified bodies.

[46] S. Gerke, T. Minssen and I. G. Cohen, "Ethical and Legal Challenges of Artificial Intelligence-Driven Healthcare" in A. Bohr and K. Memarzadeh (eds), *Artificial Intelligence in Healthcare* (Academic Press 2020) p. 295.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

post-market-surveillance machinery of the device regime can absorb the residual risk.[47] Open-ended clinical AI is different, because generative summarisation, treatment-planning support and multimodal foundation models adapted to clinical use produce an unbounded output space, and clinical ground truth is contested, so the envelope cannot be drawn tightly enough to contain the model's residual error. The device regime's intended-purpose doctrine, built for fixed-function devices, strains when a general-purpose model is brought to the bedside.[48] Healthcare is therefore a borderline case with respect to both conditions. On the supervisory side, an authority exists, but its relationship runs to the certified device rather than to the hospital's deployment of it. On the specifiability side, sufficiency tracks the task: it holds for bounded clinical AI and fails for open-ended clinical AI, which neither the Act nor existing device standards yet equip supervisors to govern.

4. *Employment: specifiability present, supervisor absent*

Employment reverses the healthcare picture and is the most instructive case. AI used in recruitment, selection, task allocation, performance evaluation and termination is high-risk under Annex III,[49] and its outputs are bounded: shortlisted or not, promoted or not, retained or dismissed. The first condition is clearly met, and existing law already addresses the right object, the decision itself. Article 22 of the GDPR governs significant automated decisions, and the Platform Work Directive adds dedicated protections,[50] while European equality law, recent scholarship stresses, is robust in substance and weak only in enforcement and transparency.[51] On our analysis the workflow is the right site of regulation, and the tools to govern it exist.

The labour-law literature confirms both halves of the picture. Abraha shows that the automated-decision provisions of the GDPR give workers real tools to contest an employment decision at the level of its output, without access to the model's internal reasoning, but that significant gaps in protection remain, above all at the point of enforcement.[52] The enforcement gap is not incidental. The features that make AI resist conventional regulatory tools in Scherer's early diagnosis, namely autonomy, foreseeability, control and opacity, weigh most heavily where the power gap between employer and worker raises the cost of a wrong decision and lowers the worker's capacity to contest it.[53] The doctrine points to the workflow; what is missing is the institution that would enforce that doctrine.

Hungarian labour-law scholarship reaches the same place from its own starting point. Bagjos and Bors treat the employment use of AI as primarily a fundamental-rights and data-protection problem, locating it in the Charter and the data-protection acquis rather than in the model.[54] Varga draws the operative consequence. Under Article 22 of the GDPR, an automated employment decision is lawful only where human involvement is genuine:

[47] AI Act, *supra*, note 3, Art. 72 (post-market monitoring); MDR, *supra*, note 44, Arts. 83–6 (post-market surveillance).

[48] Gerke, Minssen and Cohen, *supra*, note 46.

[49] AI Act, *supra*, note 3, Annex III, point 4.

[50] GDPR, *supra*, note 39, Art. 22; Directive (EU) 2024/2831 on improving working conditions in platform work [2024] OJ L, 2024/2831.

[51] A. Kelly-Lyth, "Algorithmic Discrimination at Work" (2023) 14 European Labour Law Journal 152 (existing EU equality law is "remarkably robust"); J. Adams-Prassl, R. Binns and A. Kelly-Lyth, "Directly Discriminatory Algorithms" (2023) 86 Modern Law Review 144.

[52] H. H. Abraha, "Regulating Algorithmic Employment Decisions through Data Protection Law" (2023) 14 European Labour Law Journal 172.

[53] M. U. Scherer, "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies" (2016) 29 Harvard Journal of Law & Technology 353.

[54] C. Bagjos and S. Bors, "A mesterséges intelligencia megjelenése az európai munkajogban" [The Emergence of Artificial Intelligence in European Labour Law] (2024) *Európai Jog* 2024/4, 22.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

in the words of guidance endorsed by the European Data Protection Board, the review must be "meaningful, rather than just a token gesture," and it must be performed by someone with the authority and competence to change the outcome. Nor can worker consent supply a lawful basis, because the hierarchy of the employment relationship strips it of the freedom the GDPR requires.[55] Both analyses place regulation at the decision, treat the employer as the accountable actor, and point to the ingredient we isolate, namely an institution with the standing to test whether the human in the loop is real.

That institution is absent. Employment is a standalone Annex III domain with no sectoral carve-out under Article 74, so market surveillance falls to the general Article 70 authority,[56] which in Hungary is the minister responsible for enterprise development. The ministry has no examination relationship with employers, no practice of inspecting their internal decision processes, and none of the standing the financial supervisor brings to its institutions. The data-protection authority, in Hungary the National Authority for Data Protection and Freedom of Information (NAIH), possesses the independence described above, but its practice under Article 22 is thin: its best-known AI decision, the 2022 Budapest Bank case on the emotion analysis of recorded customer-service calls, reached call-centre employees as well as customers but was decided under the transparency, lawful-basis and data-protection-by-design provisions rather than Article 22.[57] The Act adds no supervisory relationship to fill the gap. Human oversight cannot substitute for such a relationship, because oversight under Article 14 is itself a workflow-design problem: it is vulnerable to automation bias and to the reviewer reduced to a rubber stamp, and its effectiveness must itself be supervised rather than assumed.[58] The result is accountability in principle but not in practice: the framework does not fail in employment because its logic is wrong, for the object is right and the envelope is specifiable, but because the institution is missing. Employment thus isolates the second condition.

V. The redesign window: from placeholder to architecture

One statute produces three outcomes: enforceable obligations in finance, where a standing supervisor already examines the workflow; sufficiency only for bounded clinical AI in healthcare, where the supervisor's relationship stops at the artefact; and inert obligations in employment, where the workflow is specifiable but no supervisor stands behind it. The text of the Act is constant across the three; what moves the outcome is the supervisory relationship, which the Act delegates and which Hungary has, for most of the economy, declined to build. The Act's effectiveness is decided by the designation, not by the text of the Regulation.

Two developments have reopened the institutional question. The Digital Omnibus has moved the application of the high-risk obligations to December 2027 for standalone Annex

[55] D Varga, "A munkáltató és a mesterséges intelligencia" [The Employer and Artificial Intelligence] (2024) *Munkajog* 2024/3, 27; Art. 29 Working Party, "Guidelines on Automated individual decision-making and Profiling" (WP251rev.01, 6 February 2018) p. 21, endorsed by the European Data Protection Board.

[56] AI Act, *supra*, note 3, Art. 74 (no sectoral derogation for Annex III, point 4) and Art. 70.

[57] NAIH, Decision NAIH-85-3/2022 (8 February 2022), Budapest Bank Zrt., fine of HUF 250 million; on the enforcement gap generally, Abraha, *supra*, note 52, p. 172.

[58] AI Act, *supra*, note 3, Art. 14; T. Enarsson, L. Enqvist and M. Naarttijärvi, "Approaching the Human in the Loop: Legal Perspectives on Hybrid Human/Algorithmic Decision-Making in Three Contexts" (2022) 31 Information & Communications Technology Law 123; B. Green, "The Flaws of Policies Requiring Human Oversight of Government Algorithms" (2022) 45 Computer Law & Security Review 105681; R. Griffin and E. Stallman, "A Systemic Approach to Implementing the DSA's Human-in-the-Loop Requirement" (Verfassungsblog, 22 February 2024) <https://verfassungsblog.de/a-systemic-approach-to-implementing-the-dsas-human-in-the-loop-requirement/> (last accessed 29 August 2026).

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

III systems and to August 2028 for product-embedded systems.[59] Shortly before its adoption, a new Hungarian government had taken office and begun to reorganise the ministerial structure, renaming the ministry that hosts the market-surveillance department and redistributing its portfolio, so that the function's home is now in flux.[60] The placeholder of 2025 is neither settled nor fixed, and a window has opened in which to choose supervisory architecture deliberately rather than inherit it by default.

The first priority in using the window is to build the missing supervisory relationship. Horizontal AI supervision should move out of the ministry and into a dedicated authority with the independence and the examination capacity that management-based regulation requires, that is, the Hungarian Artificial Intelligence Office that the national AI Strategy first proposed.[61] Independence here is not a formality but the precondition of the examination relationship, and the standard against which to measure it already exists in the Union's law on the supervision of personal data. A department that takes its instructions from a minister would not satisfy that standard as supervisor of the data these systems process, and there is no reason why it should be taken to satisfy it as supervisor of the decisions they take. The reform is institutional, not textual: no amendment to the Regulation is needed, only a better choice under it.

The case for the ministerial solution is not frivolous: a ministry brings coordination channels across government, ready administrative resources and a clear line of political accountability, and the Act itself requires the horizontal authority to cooperate with the sectoral supervisors. These are conveniences of organisation rather than answers to the objection, because an authority that both promotes the digital economy and polices it cannot supply the impartial, examination-based relationship the technique needs. Coordination can be engineered between independent bodies; independence cannot be engineered inside a promotional ministry.

The second priority is to give the envelope operational content. The postponement gives the standardisation bodies and the AI Office time to turn the abstraction of Articles 11 and 17 into auditable workflow-governance standards that an external supervisor can verify.[62] The Act treats access to a model's source code as a last resort, which underscores that the object of verification was never the model but the governance around it, and the standards should say so.[63] A presumption of conformity is only as demanding as the harmonised standard behind it, so giving Articles 11 and 17 real content is the precondition for the conformity-assessment machinery to mean anything at the workflow level. Because the Hungarian implementing Act splits supervision between the central bank, the medical-device authority and the horizontal authority, and the AI Act requires their activities to be coordinated, the advisory council is well placed to run that coordination.[64]

The Omnibus itself complicates this prescription, because its documentation measures run the other way. Small and medium-sized providers, and now small mid-caps, may submit simplified technical documentation on a Commission form; implementation of the quality-management system is expressly proportionate to the provider's size; the template for the post-market monitoring plan has been downgraded from an implementing act to guidance due by September 2027; and the registration of systems self-assessed out of the

[59] See *supra*, note 29.

[60] Act XIII of 2026 on the list of the ministries of Hungary; Government Decree 90/2026 (V.13.) on the tasks and competences of the members of the Government; Pók, *supra*, note 33.

[61] AI Act, *supra*, note 3, Art. 70(1); Regulation (EU) 2019/1020, *supra*, note 36, Art. 11; Hungary's AI Strategy, *supra*, note 38, s 5.1.

[62] AI Act, *supra*, note 3, Arts. 11, 17, 40–1 and 96.

[63] AI Act, *supra*, note 3, Art. 74(13).

[64] AI Act, *supra*, note 3, Art. 74(10).

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

high-risk class has been streamlined.[65] Each measure is defensible as cost relief, but together they move the envelope's specification away from binding form and towards guidance, in the very instrument that concedes that the standards are late. As mandated specification recedes, more of the load shifts to the second condition. A lighter paper trail is tolerable under a supervisor who can demand the full record; without such a supervisor, it simply leaves less for anyone to examine. The simplification agenda, whatever its merits, raises the stakes of the designation choice it leaves untouched.

The diagnosis generalises beyond Hungary. Every Member State that improvised a first-wave designation faces the same question: does the body it named have, or can it build, an examination relationship with the deployers in each high-risk domain? Several have lodged horizontal supervision in data-protection authorities, which satisfy the independence condition and hold investigatory powers over every controller but have seldom examined the decision processes of employers or hospitals; others have chosen communications or cybersecurity regulators whose remit does not reach those deployers at all.[66] The text of the Act is identical in every Member State and will be applied to identical workflows; what will differ is the supervision, so the finding within Hungary, that uniform obligations produce uneven accountability wherever the supervisory relationship is uneven, holds across the Union. The window is real, and it is the last inexpensive chance to get the architecture right before the obligations apply.

VI. Conclusion

The Easterbrook–Lessig debate frames AI regulation as a binary: either general law suffices, or a technology-specific architecture is required. At the level of the model the question has no answer, because the model is neither governable by general principle nor legible as architecture. At the level of the workflow the question dissolves. Lessig is right that architecture regulates, but the architecture that matters is the deterministic envelope of the deployment workflow, which is designed, legible and auditable. Easterbrook is right that general institutions govern new technology, but only where those institutions keep a standing supervisory relationship with the deployer. The Act vindicates neither theorist: it is an instance of management-based regulation, and its success is conditional on specifiable workflows and, above all, on supervisory backing.

Hungary shows the two conditions at work within a single jurisdiction. The Act is likely to be enforced where the central bank already supervises, to work only partially where clinical tasks cannot be fully specified, and to remain a dead letter where a ministry has been designated in place of a genuine regulator. The lesson is not that the Act is too ambitious or its definitions imperfect. Rather, its effectiveness was always going to be determined by the national designation decisions rather than by the text of the Regulation itself, and the Union has now conceded as much. The postponement has given Hungary, and every Member State that improvised its first-wave designation, a brief and unrepeatable opportunity to choose its supervisory architecture deliberately, and to build the supervisory relationship that the Act presupposes but does not itself create. The priority for the next phase of AI regulation is not better definitions but better supervision. The models will keep changing; the institutions that watch over them are what will endure.

[65] Digital Omnibus, *supra*, note 7, amending AI Act, *supra*, note 3, Arts. 11(1), 17(2), 63(1) and 72(3) and Annex VIII; recitals 22, 28 and 41.

[66] European Commission, "Market Surveillance Authorities under the AI Act", *supra*, note 37.

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165

Acknowledgments. None.

Financial support. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing interests. The author declares no competing interests.

Declaration of AI use. In preparing this manuscript, the author used AI-based tools for drafting assistance, language editing and reference checking, under the author's direction and review. The author has verified all sources and content and takes full responsibility for the manuscript.

Cite this article: N Mike, "From the Law of the Horse to the Workflow: Why Institutions, Not the Model, Decide the EU AI Act's Effectiveness: Lessons from Hungary's First-Wave Implementation". *European Journal of Risk Regulation*. https://doi.org/10.1017/err.2026.10165

Downloaded from https://www.cambridge.org/core. Corvinus Egyetem, on 27 Oct 2026 at 13:55:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/err.2026.10165